UNIVERSITÀ
DEGLI STUDI
DI PADOVA

Università degli Studi di Padova

DIPARTIMENTO DI MATEMATICA "TULLIO LEVI-CIVITA"

Corso di Laurea Magistrale in Matematica

Isogeny Graphs and Cryptographic Applications

Relatore:
Dott. Rer. Nat. Alessio Caminata

Laureando: Massimo Ostuzzi
Matricola: 2055732

Correlatore:
Prof. Alberto Tonolo

Anno Accademico 2022/2023

21 Luglio 2023

Acknowledgements

I would like to spend few words for those who helped me in the realization of this work.

My special thanks go to both my supervisors: Alessio Caminata for accepting to guide me on this journey and for all the fruitful discussions; Alberto Tonolo for the help in finding a solution to write this thesis on the topic I truly wanted and for every invaluable advice.

I would like to thank my mother, my father, my sister, all my friends and volleyball teammates for supporting me throughout my whole university journey.

Massimo Ostuzzi

Contents

Introduction

Cryptography is a field of research that sits between mathematics and computer science. It is the practice and study of techniques for secure communication of confidential information. With our deepening reliance on digital communication, the need for strong cryptography becomes increasingly critical. Let us suppose that Alice and Bob want to share confidential information and that Eve wants to steal it. The aim of encryption is to enable Alice and Bob to communicate without Eve being able to recover Alice's and Bob's confidential messages. Generally speaking, there are two distinct types of encryption with two different purposes.

On the one hand, *symmetric cryptography*, also known as *secret-key cryptography*, operates under the assumption that the two communicating users have already established a common *secret key.* Symmetric cryptography is typically really efficient, and it is widely used. The most popular cryptosystems of this kind are called DES and AES.

On the other hand, *asymmetric cryptography*, also known as *public-key cryptography*, aims to allow communication without using any previously agreed secret key. This type of cryptography is mathematically deeper, and each public-key cryptosystem's security relies on some *hard* mathematical problems, which are called *primitives.* Currently, the most commonly used primitives are the *factoring problem* and the *discrete logarithm problem.*

In general, we will say that a cryptographic problem is easy if there exists a probabilistic polynomial time algorithm that solves it and we will say that it is hard if there is no known polynomial time algorithm solving it.

Asymmetric cryptography is generally much more inefficient than its symmetric counterpart. For this reason, they are usually combined: first, Alice and Bob use asymmetric cryptography to agree on a secret key, and then they use symmetric cryptography to communicate, exploiting their shared secret key.

However, there is an incumbent menace threatening secure communications: *quantum computers.* This new kind of computer can achieve unprecedented computational power, making feasible problems that were once considered computationally hard. On top of that, in 1994, Peter Shor designed a *quantum algorithm* that could potentially break all classical cryptography. A quantum algorithm is an algorithm that runs in a quantum computer. Shor's algorithm not only uses the huge computational power of quantum computers, but also exploits their quantum properties, for example *superposition* of quantum particles. Using Shor's algorithm, the factorization problem and the discrete logarithm problem become easy, i.e. solvable by a polynomial time algorithm.

This real-world problem pushed cryptographers towards the research of new cryptographic primitives with the hope of creating some new *quantum-resistant* cryptosystems.

Post-quantum cryptography is the study, implementation and analysis of cryptographic primitives and cryptosystems that are believed to be quantum-resistant. Over the past years, there have been multiple proposals in terms of new primitives. The most popular ones are based on some difficult mathematical problems in coding theory, lattice theory, multivariate systems of polynomial equations and isogeny graphs of elliptic curves.

In this thesis, our effort aims to introduce the underlying mathematics of isogeny-based cryptography. Moreover, we will present a cryptosystem called CSIDH based on isogenies, and we will analyze its functioning and properties.

More precisely, in the first chapter we will go through the theory of elliptic curves, with a special focus on finite-field elliptic curves.

In the second chapter, we will discuss the theory of elliptic curves over $\mathbb{C}$ and introduce the class group action, which is a fundamental algebraic tool to encode important geometric information.

In the third chapter, we will finally focus on the structure of isogeny graphs both for supersingular and ordinary elliptic curves, and we will prove that under some specific assumptions they are expander graphs.

In the fourth and last chapter, we will present the cryptosystem CSIDH and its features, analysing its behavior under some potential classical attacks.

We assume that the reader is familiar with the basic ideas of algebraic geometry. For a quick review of those topics, we suggest [33, Chapter 1]. If the reader wants to dive deeper into the proofs of the first three sections of the first chapter of this thesis, we suggest having a look at [33, Chapter 2], too. We also assume that the reader is no stranger to ideas from basic number theory. For example, we will need the concepts of ring of integers, number fields and Legendre symbol for inert, ramified and split primes. If the reader wants to brush up on these subjects, we suggest having a look at [26, Chapter 1] or [24].

Typically, all computations involving examples of actual elliptic curves and isogenies are really time-consuming and tedious, unless one restricts to small finite fields or some simple cases. For this reason, the ones in this thesis were (almost) all performed using the software SageMath.

Throughout this thesis, K will be a field of characteristic different from 2 and 3. This choice is convenient since, under this assumption, many proofs become easier to deal with. Moreover, the cryptosystem we will analyze is built over fields with large characteristics. Precisely, if we want to achieve quantum security in CSIDH, the characteristic of the base field needs to be greater than 2^{384}, which means it needs to have at least 116 figures.

We will call curve any irreducible one-dimensional projective variety in $\mathbb{P}^2_{\overline{K}}$.

Chapter 1

Elliptic Curves

1.1 Definitions and Group Law

In this first section, we will provide the fundamental definitions for the theory of elliptic curves. We start by recalling some basic notions of algebraic geometry for projective curves.

The *projective plane* over the field K is the quotient set

$$\mathbb{P}^2_K = \frac{K^3 \smallsetminus \{(0,0,0)\}}{\sim}$$

under the equivalence relation

$$(x, y, z) \sim (x', y', z') \quad \text{if and only if} \quad (x, y, z) = (\lambda x', \lambda y', \lambda z'),$$

for some $\lambda \in K^*$. We will denote the equivalence class of the point (x, y, z) by $[x, y, z]$. We will identify the standard affine plane over K, denoted by $\mathbb{A}^2_K$, with the subset of the projective plane

$$\mathbb{P}^2_K \cap \{[x, y, z] \in \mathbb{P}^2_K \colon z = 1\}.$$

Let C be a curve in the projective plane, $P \in C$ and $f \in \overline{K}[x, y, z]$ an homogeneous polynomial generating $I(C)$. The curve C is said *nonsingular* or *smooth* at P if we have

$$\left(\frac{\partial f}{\partial x}(P), \frac{\partial f}{\partial y}(P), \frac{\partial f}{\partial z}(P)\right) = (0, 0, 0).$$

If C is nonsingular at each point, we say that C is nonsingular.

By $I(C/K)$ we denote the ideal in $K[x, y, z]$ corresponding to the curve C/K. Given a curve C over the field K, the *function field* $K(C)$ of C/K is the field of rational functions f/g such that

(1) f and g are homogenous polynomials of the same degree in $K[x, y, z]$;

(2) $g \notin I(C/K)$;

(3) two functions f_1/g_1 and f_2/g_2 are identified if $f_1 g_2 - f_2 g_1 \in I(C/K)$.

A *rational map* of curves from C_1 to C_2 is a map of the shape

$$\phi\colon C_1 \longrightarrow C_2, \quad \phi = [f_x, f_y, f_z]$$

where the functions $f_x, f_y, f_z \in \overline{K}(C_1)$ have the property that for each point $P \in C_1$ at which f_x, f_y, f_z are all defined, we have

$$\phi(P) = [f_x(P), f_y(P), f_z(P)] \in C_2.$$

If there exists some $\lambda \in \overline{K}^*$ such that $\lambda f_x, \lambda f_y, \lambda f_z \in K(C_1)$, we say that ϕ is defined over K. The rational map ϕ is *regular* at $P \in C_1$ if there is a function $g \in \overline{K}(C_1)$ such that gf_x, gf_y and gf_z are well defined at P and at least one of them is nonzero in P. If such g exists, we set

$$\phi(P) = [gf_x(P), gf_y(P), gf_z(P)].$$

Notice that it may be necessary to consider different g for different points. A rational map that is regular at each point is called *morphism* of curves. We will often say *K-rational* instead of defined over K.

One can easily transfer the above definitions to the affine case. Indeed, most of the time, we will work with affine coordinates.

Definition 1.1.1. Let E be the zero locus in $\mathbb{P}^2_{\overline{K}}$ of a cubic equation with only one point on the line at infinity, called the base point. If E is nonsingular, then we say it is an *elliptic curve.*

Throughout the whole dissertation, we will only consider fields of characteristic different from 2 and 3. Under this hypothesis, any elliptic curve E can be described as the zero locus of a *short Weierstrass equation*

$$y^2 = x^3 + Ax + B$$

with $A, B \in \overline{K}$. Notice that its projective closure only has one point at infinity, namely the point $O = [0, 1, 0]$. If $A, B \in K$, we say that the curve is defined over K and we will write E/K. If $P \in V$ has coordinates in K, we say it is a *K-rational point.*

Clearly, not all such short Weierstrass equations correspond to an elliptic curve, since we must have smoothness. Applying the condition for nonsingularity to the projectivization of the short Weierstrass equation

$$y^2z = x^3 + Axz^2 + Bz^3,$$

we get the system of equations

$$\begin{cases} 3x^2 + Az^2 = 0 \\ 2yz = 0 \\ y^2 = 2Axz + 3Bz^2. \end{cases}$$

From the second equation, we have $y = 0$ or $z = 0$. If $z = 0$, then we must have $x = 0$ and $y = 0$, which does not yield a projective point. Hence, we can assume $z \neq 0$ and set $z = 1$ to reduce to the affine case. We get a new system of equations

$$\begin{cases} 3x^2 + A = 0 \\ y = 0 \\ y^2 = 2Ax + 3B, \end{cases}$$

and exploiting the second equation we get

$$\begin{cases} 3x^2 + A = 0 \\ 2Ax + 3B = 0. \end{cases}$$

If $A = 0$, then we also have $B = 0$ and so the curve of equation

$$y^2 = x^3$$

has a singular point in $(0, 0)$. This means that this curve is not an elliptic curve. Assume now that $A \neq 0$. Then, combining the two equations of the system, we are left with the relation

$$4A^3 + 27B^2 = 0.$$

Thus, if we want to rule out singular curves, we can require that the last equation is not satisfied. Notice that it also works in the case $A = 0$.

The previous computations yield the following definition and proposition.

Definition 1.1.2. Let E be the zero locus of a short Weierstrass equation

$$y^2 = x^3 + Ax + B.$$

We define the *discriminant* of E to be

$$\Delta(E) = -16(4A^3 + 27B^2).$$

Proposition 1.1.3. *A curve E given in short Weierstrass equation is an elliptic curve if and only if $\Delta(E) \neq 0$.*

Another fundamental quantity in the theory of elliptic curves is the j-invariant.

Definition 1.1.4. Let E be an elliptic curve given in short Weierstrass equation

$$y^2 = x^3 + Ax + B.$$

We define the *j-invariant* to be

$$j(E) = -1728\frac{(4A)^3}{\Delta}.$$

Notice it is well defined only for elliptic curves, since their discriminant is nonzero.

Proposition 1.1.5. *Two elliptic curves are isomorphic over $\overline{K}$ if and only if they have the same j-invariant. Moreover, let $j_0 \in \overline{K}$. There exists an elliptic curve defined over $K(j_0)$ whose j-invariant equals j_0.*

Proof. For a proof of the first statement, see [33, Proposition III.1.4]. Let $j_0 \neq 0, 1728$ and consider the equation

$$y^2 + xy = x^3 + \frac{36}{j_0 - 1728}x + \frac{1}{j_0 - 1728}. \tag{1.1}$$

Performing a coordinate change, we can put it in the short Weierstrass form

$$y^2 = x^3 - \frac{j_0}{12\sqrt[3]{4}(j_0 - 1728)}x + \frac{j_0}{216(j_0 - 1728)}.$$

Simple calculations show it is a nonsingular curve and its j-invariant is j_0. Hence it is the desired elliptic curve. Notice that from the short Weierstrass form it is not clear that the curve is defined over $K(j_0)$, since we cannot assume that $\sqrt[3]{4} \in K(j_0)$. Notice that 12 and 216 are surely nonzero in K, as we assumed the characteristic of K is different from 2 and 3. However, looking at the Equation (1.1) which is not in short Weierstrass form, we can see that it is indeed defined over $K(j_0)$. To complete the picture, we provide the two curves

$$y^2 = x^3 + 1, \tag{1.2}$$
$$y^2 = x^3 + x. \tag{1.3}$$

The curve given by Equation (1.2) has discriminant and j-invariant $\Delta = -27$ and $j = 0$ and it is clearly defined over $K(0) = K$. The curve given by Equation (1.3) has discriminant and j-invariant $\Delta = -64$ and $j = 1728$ and it is clearly defined over $K(1728) = K$, too. □

In general, if two curves are isomorphic over the algebraic closure of the base field, we will simply say that they are isomorphic.

Corollary 1.1.6. *An elliptic curve E is defined over K if and only if $j(E) \in K$.*

Proof. This follows directly from Proposition 1.1.5. □

Corollary 1.1.7. *Let E be an elliptic curve in short Weierstrass equation*

$$y^2 = x^3 + Ax + B,$$

any other short Weierstrass equation for E is of the form

$$y^2 = x^3 + u^4Ax + u^6B,$$

for some $u \in \overline{K}^*$, and the change of coordinates is given by the map

$$(x, y) \longmapsto (u^2x, u^3y).$$

Proof. Let

$$y^2 = x^3 + A'x + B'$$

be another short Weierstrass equation for E. Clearly, the two short Weierstrass equations for E must yield two isomorphic curves. Hence, their j-invariants coincide, i.e. we have

$$\frac{4A^3}{4A^3 + 27B^2} = \frac{4A'^3}{4A'^3 + 27B'^2}.$$

This yields a relation between the coefficients

$$A^3B'^2 = A'^3B^2.$$

If $A = 0$ or $B = 0$ in K, then we must have $A' = 0$ or $B' = 0$, respectively. Hence, if one of the two coefficients is zero, all possible values of the other one yield the same curve up to isomorphism. Notice that we cannot have both coefficients simultaneously zero, as E is nonsingular by hypothesis.

Suppose now that both A and B are nonzero in K. Let us denote $u = \sqrt[6]{B'/B}$. We can directly deduce the relations

$$A' = u^4A \quad \text{and} \quad B' = u^6B,$$

which show our initial claim. □

Example 1.1.8. Let E_1, E_2 and E_3 be the curves of short Weierstrass equations

$$y^2 = x^3 + x, \quad y^2 = x^3 + 12x + 12 \quad \text{and} \quad y^2 = x^3 - 3x + 2$$

over $\mathbb{F}_{19}$, the finite field with 19 elements. We can directly compute the discriminants, which are $\Delta(E_1) \equiv 12 \mod 19$, $\Delta(E_2) \equiv 5 \mod 19$ and $\Delta(E_3) \equiv 0 \mod 19$, and conclude that only E_1 and E_2 are actually elliptic curves. Their j-invariants are 18 and 10, respectively, which are not congruent modulo 19. Thus, they are not isomorphic. In Figure 1.1 we can visualize their $\mathbb{F}_{19}$-rational affine points.

Example 1.1.9. Let us consider the two elliptic curves F_1 and F_2 over $\mathbb{R}$ with equations

$$y^2 = x^3 - 4x + 5 \quad \text{and} \quad y^2 = x^3 - 2x + 1,$$

respectively. In Figure 1.2, we can observe a graphical representation of their affine points. The only missing point is the base point $O = [0, 1, 0]$, which is not in the picture as it lies in the line at infinity.

Remark 1.1.10. A priori, two elliptic curves defined over $K \neq \overline{K}$ can be isomorphic over $\overline{K}$ but not over K. For example, if ϕ is an isomorphism defined over a degree two extension of K but not over K, we call ϕ a quadratic twist. If the extension has degree three, four or six, we call ϕ a cubic, quartic or sextic twist, respectively. Every E/K

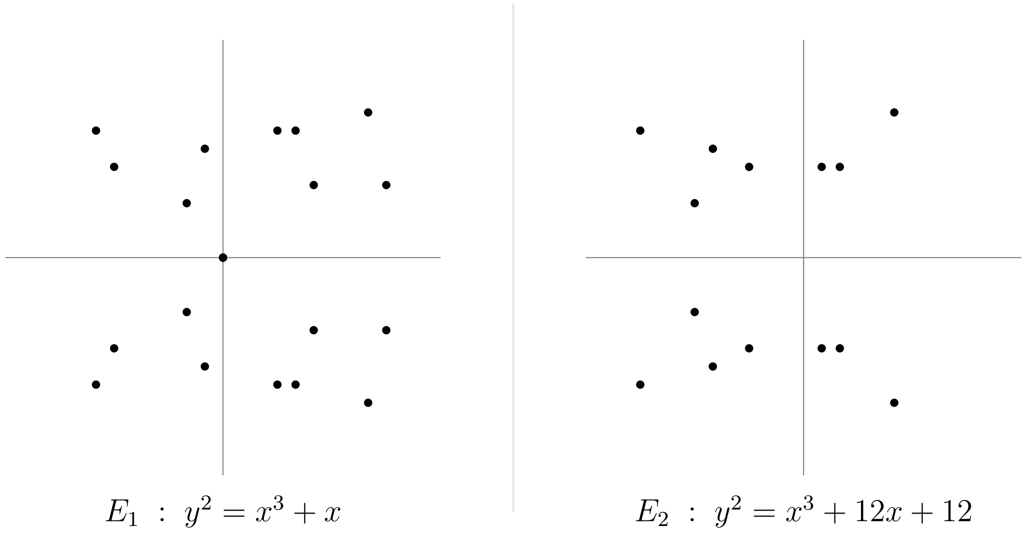

$E_1 \ : \ y^2 = x^3 + x$ $\qquad$ $E_2 \ : \ y^2 = x^3 + 12x + 12$

Figure 1.1: A graphical visualization of the two elliptic curves E_1 and E_2 over $\mathbb{F}_{19}$ of Example 1.2.9.

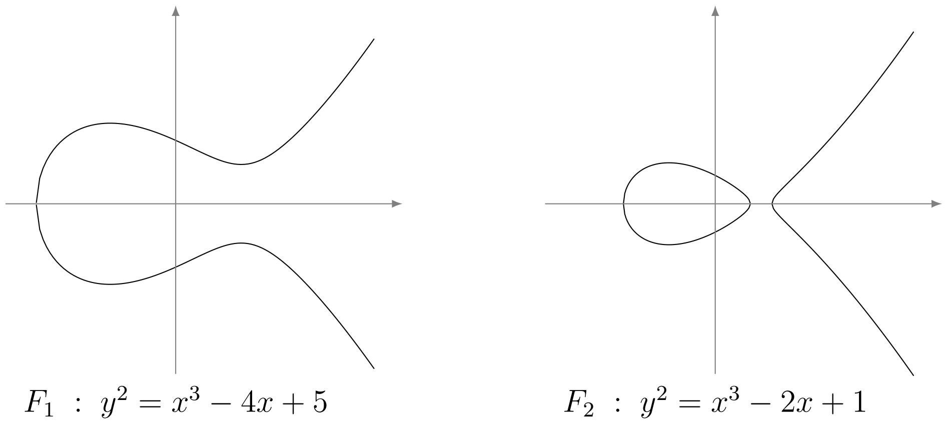

$F_1 \ : \ y^2 = x^3 - 4x + 5$ $\qquad$ $F_2 \ : \ y^2 = x^3 - 2x + 1$

Figure 1.2: Affine points of the real elliptic curves F_1 and F_2 from Example 1.1.9.

elliptic curve admits a quadratic twist. Let us suppose there exists $u \in K$ such that u is not a square in K. Then, if E has short Weierstrass equation

$$y^2 = x^3 + Ax + B,$$

we can define E' to be the elliptic curve defined over K of equation

$$uy^2 = x^3 + Ax + B.$$

We can put E' in short Weierstrass equation substituting x/u for x and y/u^2 for y and we get

$$y^2 = x^3 + u^2Ax + u^3B.$$

The curves E and E' are isomorphic over $K(\sqrt{u})$ but not over K, since u is not a square in K. The isomorphism is

$$\phi\colon E \longrightarrow E', \quad \phi(x, y) = (x, y\sqrt{u}).$$

The two curves of j-invariant $j = 0$ and $j = 1728$ also have cubic, quartic and sextic twists.

At this point, we are ready to endow elliptic curves with an algebraic structure. Let E be an elliptic curve in short Weierstrass equation, $P, Q \in E$ and L be the line in $\mathbb{P}^2_K$ through P and Q. If $P = Q$, take L as the tangent line to E at P. Since the elliptic curve has degree three, by Bezout theorem, we have that the line L intersects E in a third point R, possibly P or Q. Let L' be the line in $\mathbb{P}^2$ through R and O, the base point of the curve. Again, this line intersects E in a third point, possibly coinciding with R or O, which we will denote by $P + Q$. This is a well defined operation that produces a third point, namely $P + Q$, starting from any two points P and Q on the curve.

Proposition 1.1.11. *With the above notations, the operation $+\colon E \times E \longrightarrow E$ satisfies the following properties.*

(1) *It is commutative.*

(2) *$P + O = O$ for all $P \in E$, i.e. the point O acts as the identity element of the group.*

(3) *Every point has an inverse, which means that for each $P \in E$ there exists $-P \in E$ such that $P + (-P) = O$.*

(4) *It is associative.*

(5) *If a line intersects E in the three points P, Q, R, then $P + Q + R = O$.*

Thus, the operation $+$ makes E into an abelian group with identity element O. Moreover, if E is defined over K, then

$$E(K) = \{(x, y) \in K^2 \colon y^2 = x^3 + Ax^2 + B\} \cup \{O\}$$

is a subgroup of E. We will write simply E for the set of points $E(\overline{K})$.

Proof. All claims except the fourth one are trivial and can be deduced directly from the construction of the group law. See [33, Proposition III.2.2]. □

One can deduce explicit formulas for the addition and negation by working with the coordinates and the Weiestrass equation.

Proposition 1.1.12. *Let E/K be an elliptic curve of short Weierstrass equation*

$$y^2 = x^3 + Ax + B.$$

Let $P = (x_P, y_P)$ be a point in E. We have $-P = (x_P, -y_P)$. Now, let $P_i = (x_i, y_i) \in E$ for $i = 1, 2, 3$ such that $P_1 + P_2 = P_3$.

(1) *If $x_1 \neq x_2$, we have*

$$x_3 = \lambda^2 - x_1 - x_2, \quad y_3 = \lambda(x_1 - x_3) - y_1$$

where $\lambda = \dfrac{y_2 - y_1}{x_2 - x_1}$.

(2) *If $P_1 = P_2$ and $y_1 \neq 0$, we have*

$$x_3 = \nu^2 - 2x_1, \quad y_3 = \nu(x_1 - x_3) - y_1,$$

where $\nu = \dfrac{3x_1^2 + A}{2y_1}$.

Proof. See [33, III.2.3] for a detailed proof. □

Example 1.1.13. Let us consider again the curves E_1 and E_2 from Example 1.1.8, whose equations are

$$y^2 = x^3 + x, \quad \text{and} \quad y^2 = x^3 + 12x + 12.$$

respectively. Notice that the point $P = (0, 0) \in E_1$ is not the identity element of the group and has clearly order 2. The point $Q = (1, 5) \in E_2$ has not order two. Indeed, following the notation in Proposition 1.1.12, we set $P_1 = P_2 = Q$, so that

$$\nu = \frac{15}{10} = 11$$

in $\mathbb{F}_{19}$. Plugging this value into the formula, we get

$$x_3 = 121 - 2 \equiv 5 \mod 19 \quad \text{and} \quad y_3 = 11(1 - 5) - 5 \equiv 8 \mod 19.$$

Hence we have $[2]Q = (5, 8) \in E_2$.

Using SageMath, one can verify that the point $[1, 5, 1] \in E_2$ has order 15 in the group.

Now, we give the second definition of elliptic curves, which is more intrinsic, since it is coordinate-free.

The following theorem shows that the addition and negation maps are actually morphisms, in the sense that they are rational maps defined at each point of the curve.

Theorem 1.1.14. *Let E/K be an elliptic curve. Then the equations giving the group law coordinate-wise on E define morphisms*

$$+ \colon E \times E \longrightarrow E \qquad - \colon E \longrightarrow E.$$

Proof. For a complete and precise proof, see [33, Theorem III.3.6]. Here, we just give an idea of the actual proof, without diving into the precise computations with coordinates. Looking at the explicit formulas, one can conclude immediately that the negation map is a morphism: the formula is given by a quotient of polynomials and the source of the morphism is smooth, so the map is regular at each point, by Proposition 1.2.1. For the addition map, we cannot directly use this strategy. Indeed, it is a rational map whose source is $E \times E$, a variety of dimension two. The trick consists of fixing a point P in E and considering the translation map by P. Now, looking at explicit formulas, this is a morphism by the same argument we used for the negation map. Continuing in this direction and working out some particular cases separately, one can conclude the proof. □

1.2 Isogenies

In this chapter, we want to focus on a particular class of morphisms of elliptic curves.

We begin with two classical results for morphisms of general curves.

Proposition 1.2.1. *Let C be a curve, $V \subseteq \mathbb{P}^N$ a variety and $P \in C$ a smooth point. Let $\phi\colon C \longrightarrow V$ be a rational map. Then ϕ is regular at P.*

Proof. See [33, Proposition II.2.1]. □

This means in particular that every rational map between elliptic curves is a morphisms, i.e. it's regular at each point.

Proposition 1.2.2. *Let $\phi\colon C_1 \longrightarrow C_2$ be a morphism of curves. Then ϕ is either constant or surjective.*

Proof. The image of ϕ must be irreducible, since C_1 is irreducible. Hence it can be either the whole C_2 or a single point. □

Definition 1.2.3. Let E_1 and E_2 be elliptic curves in short Weierstrass equation. An *isogeny* from E_1 to E_2 is a morphism of elliptic curves

$$\phi\colon E_1 \longrightarrow E_2 \quad \text{satisfying} \quad \phi(O) = O.$$

We say that two elliptic curves E_1/K and E_2/K are *isogenous* over K if there exists a surjective K-rational isogeny from E_1 to E_2.

Remark 1.2.4. Since every elliptic curve is smooth, an isogeny can be either the constant map $\phi(E_1) = O$ or surjective.

Every isogeny, except for the constant one, is a *finite map* of curves, which means that the preimage of each point on the target has finite cardinality. See [16, Proposition II.6.8]. We define a map associated with each isogeny ϕ, called *pull-back* of ϕ.

Definition 1.2.5. Let E_1 and E_2 be elliptic curves defined over K, $\phi\colon E_1 \longrightarrow E_2$ a K-rational isogeny. We define the pull-back as the map

$$\phi^*\colon K(E_2) \longrightarrow K(E_1)$$

such that $\phi^*(f) = f \circ \phi$ for each $f \in K(E_2)$. The degree of the isogeny ϕ, denoted by $\deg \phi$, is the degree of the finite extension of fields $K(E_1)/\phi^* K(E_2)$ and similarly for the inseparable and separable degrees, denoted by $\deg_i \phi$ and $\deg_s \phi$ respectively. We say that the map ϕ is *separable*, *inseparable* or *purely inseparable* according to the corresponding property of the field extension. By convention, we fix $\deg[0] = 0$, where $[0]$ is the constant isogeny.

Remark 1.2.6. The degree is multiplicative with respect to the composition of isogenies. Moreover, an isogeny of degree one is indeed an isomorphism. This follows from the fact that there is an underlying contravariant equivalence of categories. Consider the following two categories:

(1) *K-Curves*, whose objects are smooth curves defined over K and morphisms are nonconstant morphisms defined over K of smooth curves.

(2) *Ext-K*, whose objects are finite field extensions $\mathbb{K}/K$ of transcendence degree one and morphisms are field injections fixing K.

We define the contravariant functor

$$^*\colon \textit{K-Curves} \longrightarrow \textit{Ext-K}$$

that at the level of objects associates each curve C to its function field $K(C)$, while at the level of morphisms associates each $\phi\colon C_1 \longrightarrow C_2$ morphism of smooth curves to its pullback ϕ^*. This functor encodes the underlying connection between algebra and geometry in this situation. Indeed, it is actually a contravariant equivalence of categories.

Isogenies are preferable to general morphisms because they respect the abelian group structure, too. Indeed, the following result holds.

Proposition 1.2.7. *Let $\phi\colon E_1 \longrightarrow E_2$ be an isogeny between two elliptic curves. Then $\phi(P+Q) = \phi(P) + \phi(Q)$ for each $P, Q \in E_1$.*

Proof. See [33, Proposition III.4.8]. □

Corollary 1.2.8. *Let $\phi\colon E_1 \longrightarrow E_2$ be a nonconstant isogeny of elliptic curves. Then*

$$\ker \phi = \phi^{-1}(O)$$

is a finite subgroup.

Proof. This follows from Proposition 1.2.7 and the fact that isogenies are finite morphisms. □

An isogeny of degree n is customarily called *n-isogeny.* If the kernel of an isogeny ϕ is a cyclic group, we say that ϕ is a *cyclic isogeny.*

Since elliptic curves are abelian groups, morphisms between them form groups. We denote isogenies from E_1 to E_2 by $\mathrm{Hom}_K(E_1, E_2)$, and if the field is not specified we mean isogenies over $\overline{K}$. The sum of two isogenies is defined by $(\phi + \psi)(P) = \phi(P) + \psi(P)$ and Theorem 1.1.14 implies that this is an isogeny. If $E_1 = E_2$, we can also compose isogenies. Then we denote by $\mathrm{End}_K(E) = \mathrm{Hom}_K(E, E)$ the ring whose addition law is given by the sum of isogenies and whose multiplication is the composition. We call it *endomorphism ring of E.* We always follow the convention that, if the field is not specified, we mean we are considering the endomorphism ring over $\overline{K}$.

Example 1.2.9. Let $E/\mathbb{F}_{19}$ and $F/\mathbb{F}_{19}$ be two elliptic curves whose Weierstrass equations are

$$y^2 = x^3 + x, \quad \text{and} \quad y^2 = x^3 + 15x.$$

Using the software SageMath, we can find the rational map

$$\phi\colon E \longrightarrow F, \quad (x, y) \longmapsto \left(\frac{x^2+1}{x}, \frac{x^2y - y}{x^2} \right).$$

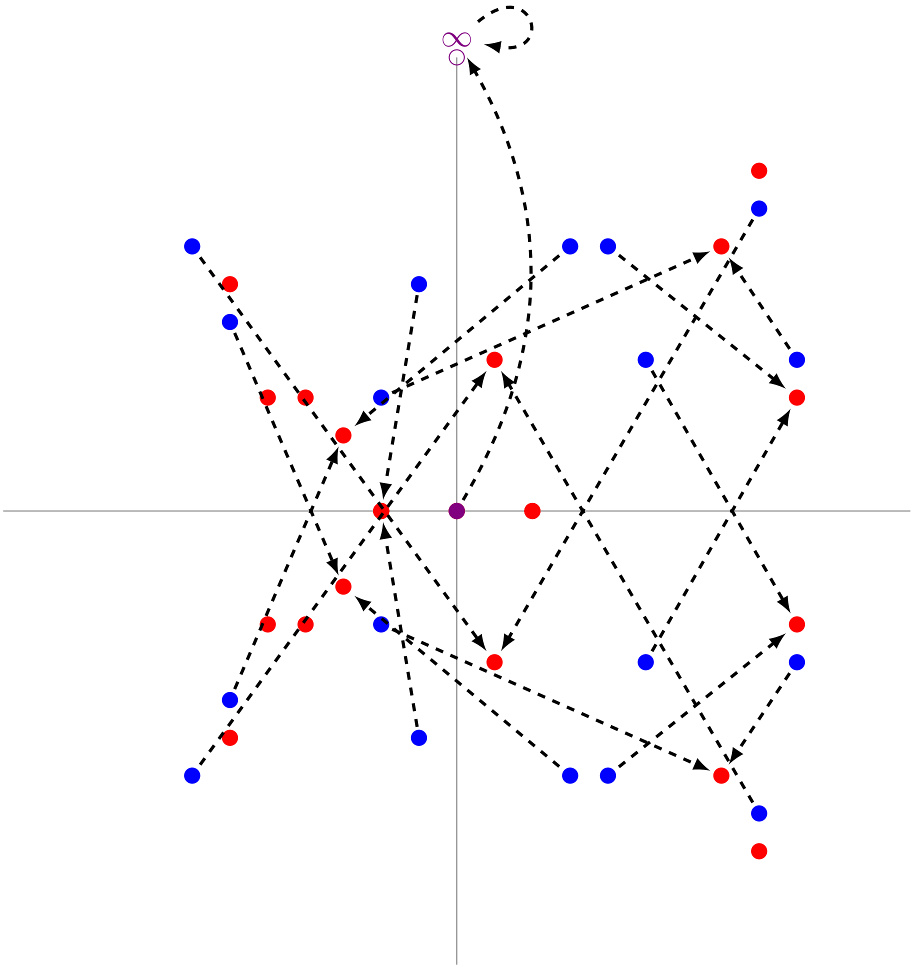

Figure 1.3: The isogeny ϕ of Example 1.2.9, as a map between curves defined over $\mathbb{F}_{19}$. The blue dots represent E, while the red ones represent F. The point of coordinate $(0, 0)$ and the empty dot on top representing the point at infinity $O = [0, 1, 0]$ are violet, since they lie both in E and F.

This is a nonconstant rational map between elliptic curves defined over $\mathbb{F}_{19}$, so it is surjective. We have $\phi(O) = O$. Hence ϕ is an isogeny.

Using again SageMath, one can verify the following claims. ϕ has degree 2 and its kernel is the subgroup of order 2 generated by the affine point $(0, 0) \in E$. In the following lines, we test with an example that this isogeny is actually an homomorphism of groups, as we expect from Proposition 1.2.7. Let us consider the points $P = (5, -4)$ and $Q = (3, 7)$. A direct computation shows that both points are in E. Plugging the coordinates of P and Q in the formula for ϕ, one can verify that

$$\phi(P) = (9, 3) \quad \text{and} \quad \phi(Q) = (-3, 2).$$

Performing the addition on F, we have $\phi(P) + \phi(Q) = (1, 4)$. On the other hand, we have $\phi(P + Q) = \phi(8, -8) = (1, 4)$, as we wanted.

In Figure 1.3, we can observe a graphical representation of the map ϕ. As a visual proof of the fact that ϕ is also a group morphism, we can observe that in the picture there is a symmetry with respect to the x-axis, the reason being that for each $P \in E$ we have $\phi(-P) = -\phi(P)$.

One may be surprised that the map is not surjective, as predicted by Proposition 1.2.2. For example, the point $(0,0) \in F$ is not in the image of ϕ. However, the map needs to be surjective if we consider not only $\mathbb{F}_{19}$-rational points, but all points of E and F over the algebraic closure $\overline{\mathbb{F}}_{19}$.

Theorem 1.2.10. *Let $\phi\colon E_1 \longrightarrow E_2$ be a nonconstant isogeny. For each $Q \in E_2$, it holds*

$$\#\phi^{-1}(Q) = \deg_s \phi.$$

Proof. For any type of curve and any morphism of curves, we have $\#\phi^{-1}(Q) = \deg_s \phi$ for all but finitely many points. If we work with elliptic curves and isogenies, for each $Q, Q' \in E_2$, we can find a point $T \in E_1$ such that $\phi(T) = Q - Q'$, as the isogeny is nonconstant, hence surjective. Since ϕ is an homomorphism of abelian groups, there is a on-to-one correspondence

$$\phi^{-1}(Q) \longrightarrow \phi^{-1}(Q'), \qquad P \longmapsto P + T,$$

which proves the claim. □

Corollary 1.2.11. *Every purely inseparable isogeny has trivial kernel.*

Proof. If an isogeny is purely inseparable, then its separable degree equals one. Thus, by the previous theorem, we have that $\#\phi^{-1}(O) = 1$. □

For each $m \in \mathbb{Z}$, we can define the *m-multiplication isogeny*

$$[m]\colon E \longrightarrow E$$

such that

$$[m](P) = \underbrace{P + ... + P}_{m}$$

if $m > 0$, and $[m](P) = [-m](-P)$ if $m < 0$. This is clearly an isogeny, because it fixes O and, looking at explicit formulas, one can check that it is a morphism. Furthermore, if E is defined over K, then $[m]$ is defined over K.

Proposition 1.2.12. *Let E/K be an elliptic curve and let $m \in \mathbb{Z}$ with $m \neq 0$. Then the m-multiplication is a nonconstant isogeny.*

Proof. See [33, Proposition III.4.2].

□

Lemma 1.2.13. *Let E/K be an elliptic curve and let $m \in \mathbb{Z}$. Assume that $m \neq 0$ in K. Then the m-multiplication map is a separable isogeny.*

Proof. See [33, Corollary III.5.4]. □

Remark 1.2.14. We can embed $\mathbb{Z}$ into $\mathrm{End}_K(E)$ for each E/K elliptic curve:

$$\mathbb{Z} \longrightarrow \mathrm{End}_K(E), \quad m \longmapsto [m].$$

Indeed, if E is defined over K, then also $[m]$ is defined over K for each $m \in \mathbb{Z}$. Sometimes we will write simply m instead of $[m]$, if the context is clear. Notice that the maps $[m]$ commute with each map in End(E), for every $m \in \mathbb{Z}$. This is because isogenies are also group morphisms.

Corollary 1.2.15. *The group of isogenies between two elliptic curves is a torsion-free $\mathbb{Z}$-module. Moreover, the endomorphism ring of an elliptic curve is a characteristic 0 integral domain.*

Proof. Given $\phi \in \mathrm{Hom}(E_1, E_2)$ and $m \in \mathbb{Z}$, if

$$[m] \circ \phi = [0],$$

then we must have $(\deg[m])(\deg \phi) = 0$. Thus, by the previous proposition, either $m = 0$ or $\phi = [0]$. This proves the first statement. Moreover, this proves also that the endomorphism ring of an elliptic curve has characteristic 0. To conclude the proof, given an elliptic curve E, suppose that $\phi, \psi \in \mathrm{End}(E)$ satisfy $\phi \circ \psi = [0]$. Then taking their degrees, we observe that at least one of them must have degree zero, which implies that it is the constant isogeny. Hence $\mathrm{End}(E)$ is an integral domain.

□

Next, we build a very special isogeny in positive characteristic, which is fundamental in the theory of elliptic curves.

Definition 1.2.16. Let K be a field with positive characteristic p and $q = p^r$. Let E/K be an elliptic curve given by a short Weierstrass equation

$$y^2 = x^3 + Ax + B.$$

We define the curve $E^{(q)}$, for $r > 0$, to be the zero locus of the equation

$$y^2 = x^3 + A^q x + B^q.$$

We have a canonical morphism $\pi_q \colon E \longrightarrow E^{(q)}$ such that, for each $(x, y) \in E$,

$$\pi_q(x, y) = (x^q, y^q).$$

This is called *q-Frobenius morphsim.*

Notice that

$$\Delta(E^{(q)}) = \Delta(E)^q \quad j(E^{(q)}) = j(E)^q.$$

Thus, $E^{(q)}$ is an elliptic curve and π_q is an isogeny, as it fixes the point at infinity.

Lemma 1.2.17. *Let K be a field of characteristic p and $q = p^r$ for some r natural number. Let E/K be an elliptic curve. Let π_q be the q-Frobenius morphism of E. The following are true:*

1. *π_q is a purely inseparable isogeny;*

2. *π_q has degree q.*

Proof. See [33, Proposition II.2.11]. □

Remark 1.2.18. Let us assume that $K = \mathbb{F}_q$. Then, the q-Frobenius morphism becomes an endomorphism, since $E = E^{(q)}$. Moreover, notice that it acts as the identity on the points of $E(\mathbb{F}_q)$. More precisely, the set of fixed points of the Frobenius endomorphism is exactly the finite group $E(\mathbb{F}_q)$. Furthermore, note that the q-Frobenius endomorphism commutes with every element of $\text{End}(E)$. This follows from the fact that for any rational function $f \in \mathbb{F}_q(x_0, ..., x_n)$ we have $f(x_0, ..., x_n)^q = f(x_0^q, ..., x_n^q)$. Thus, the subring $\mathbb{Z}[\pi_q]$ lies in the center of $\text{End}(E)$.

Next we introduce the following lemma, which claims that, in positive characteristic fields, each isogeny can be written as a composition of a separable isogeny with the p^r-Frobenius morphism, for some $r > 0$.

Lemma 1.2.19. *Every isogeny $\phi\colon E_1 \longrightarrow E_2$ over a field K of positive characteristic factors as a composition $\phi = \psi \circ \pi_q$, where $q = deg_i \phi$ and the isogeny ψ is separable.*

Proof. See [33, Corollary II.2.12]. □

Remark 1.2.20. If we work in positive characteristic, this lemma tells us that separable isogenies are the most interesting ones, since the inseparable part can be always written using Frobenius morpshisms. Actually, in characteristic 0 all isogenies are separable, since every field extension of characteristic 0 is separable.

Using this lemma, we can prove an important result that lies at the heart of Hasse estimate.

Proposition 1.2.21. *Let K be a field with positive characteristic p. Let π be the Frobenius endomorphism. Then the isogeny $\pi - [1]$ is separable.*

Proof. First, we prove that the sum of two inseparable morphism is inseparable. Let $\phi, \psi\colon E_1 \longrightarrow E_2$ be inseparable isogenies. Then, by the above lemma, there exist ϕ', ψ' separable isogenies and $r, s \geq 0$ such that

$$\phi = \phi' \circ \pi_{p^r}, \quad \psi = \psi' \circ \pi_{p^s}.$$

Actually, both being inseparable, r and s must be strictly positive. Thus, we have

$$\phi + \psi = \phi' \circ \pi_{p^r} + \psi' \circ \pi_{p^s} = (\phi' \circ \pi_{p^{r-1}} + \psi' \circ \pi_{p^{s-1}}) \circ \pi_p,$$

which is inseparable because it p divides its inseparable degree. Now we can prove our claim. Notice that the isogeny $[1]$ is separable, since it is the identity map. Assume by contradiction that $\pi - [1]$ is inseparable. Then so is $-(\pi - [1])$, and $[1] = \pi - (\pi - [1])$ is inseparable as it is written as the sum of two inseparable maps. □

The next result is arguably the most important theorem on isogenies. It allows to identify isogenies with their kernels.

Theorem 1.2.22. *Let E be an elliptic curve and G be a finite subgroup of E. There are an elliptic curve E' and a separable isogeny $\phi\colon E \longrightarrow E'$ unique up to isomorphism such that*

$$\ker \phi = G.$$

Proof. Here we try to sum up the idea of the proof. We can associate with each point P of the curve E a morphism τ_P, the translation map, which is invertible but not an isogeny, if $P \neq O$. In this way, we can associate a group of automorphism $\mathcal{G}$ to each finite subgroup G of E, consisting of translation maps. Thus, we obtain a morphism $\phi\colon E \longrightarrow E/\mathcal{G}$, where the quotient is taken according to the action of $\mathcal{G}$ on E. One can prove that the quotient is still an elliptic curve, using Hurwitz Theorem to compute its genus. One should also prove that ϕ is separable and verify that the kernel is G. For a more precise proof, see [33, Proposition III.4.12]. □

Remark 1.2.23. It is customary to denote the curve E' with the notation E/G. In the following, we will introduce two theorems due to Velú that provide an explicit expression for the isogeny ϕ and the curve E' of the previous theorem. From these, one can immediately see that if G is $\mathrm{Gal}(\overline{K}/K)$-invariant, then ϕ is defined over K. This observation will be useful in Chapter 4.

Corollary 1.2.24. *An isogeny of composite degree can always be decomposed into a sequence of isogenies of prime degrees.*

Proof. Let $\phi\colon E_1 \longrightarrow E_2$ be an isogeny of elliptic curves over K. Without loss of generality, we can assume that ϕ is separable. Indeed, if K has characteristic 0, then it is separable, since every field extension of characteristic 0 is separable. Otherwise, if K has positive characteristic p, we can reduce to consider the separable part, using Lemma 1.2.19 and the fact that the p-Frobenius morphism has degree p, which is prime. As a non-trivial abelian group, $G = \ker \phi$ contains a subgroup H of prime order. By Theorem 1.2.22, there exists a separable isogeny $\phi_1\colon E_1 \longrightarrow E_3$ with kernel H. Notice that ϕ_1 has prime degree, equal to the order of H. Then $\phi_1(G)$ is a finite subgroup of E_3 isomorphic to G/H and applying Theorem 1.2.22 again we can find a separable isogeny $\phi_2\colon E_3 \longrightarrow E_4$ with kernel $\phi_1(G)$. The kernel of the composition $\phi_2 \circ \phi_1$ is precisely G, by construction. Hence, by uniqueness in Theorem 1.2.22, there is an isomorphism $\iota\colon E_4 \longrightarrow E_2$ such that $\phi = \iota \circ \phi_2 \circ \phi_1$. Now one can proceed by induction and apply the same procedure to the isogeny $\iota \circ \phi_2$, which has smaller degree than ϕ. Eventually, one obtains a sequence of separable isogenies of prime degree whose composition is equal to ϕ. □

Starting from any finite subgroup G of E, we can explicitly compute equations for the unique isogeny with kernel G and the unique target elliptic curve, using Velú formulas. Let E be an elliptic curve of equation $y^2 = x^3 + Ax + B$ and let G be a finite subgroup of E. Let G_0 be the set of nonzero points in G. They all are affine points $Q = (x_Q, y_Q)$.

For each $P = (x, y) \in E \smallsetminus G$ let us define

$$\phi(P) = \left(x + \sum_{Q \in G_0} (x_{P+Q} - x_Q), y + \sum_{Q \in G_0} (y_{P+Q} - y_Q) \right),$$

where x_{P+Q} and x_{P+Q} are the affine coordinates of the point $P+Q$, which we can view as rational functions of x, y, x_Q, y_Q. This function clearly defines a rational map, so it defines a morphism from E to some smooth projective curve E'. Moreover, the group law of E induces a group law on E' that is defined by rational maps, hence E' is an elliptic curve too. Furthermore, for any $P \in E \smallsetminus G$ we have $\phi(P) = \phi(P + Q)$ if and only if $Q \in G$, so $\ker \phi = G$. Thus, if ϕ is separable, it is the isogeny we are looking for. By using the group law, one can get explicit formulas both for ϕ and E'. The details for these computations are convoluted. See [36, Theorem 12.16].

Theorem 1.2.25 (Velú)**.** *Let E be an elliptic curve over K of equation $y^2 = x^3 + Ax + B$ and let $x_0 \in \overline{K}$ be a root of $x^3 + Ax + B$. Define $t = 3x_0^2 + A$ and $w = x_0 t$. The rational map*

$$\phi(x, y) = \left(\frac{x^2 - x_0 x + t}{x - x_0}, \frac{(x - x_0)^2 - t}{(x - x_0)^2} y \right)$$

is a separable isogeny from E to E' with equation $y^2 = x^3 + (A - 5t)x + B - 7w$. The kernel of ϕ is the group of order 2 generated by $(x_0, 0)$.

Theorem 1.2.26 (Velú)**.** *Let E be an elliptic curve over K of equation $y^2 = x^3 + Ax + B$ and let G be a finite subgroup of E of odd order. For each nonzero $Q = (x_Q, y_Q) \in G$, let $t_Q = 3x_Q^2 + A$, $u_Q = 2y_Q^2$ and $w = u_Q + t_Q x_Q$. Let*

$$t = \sum_{Q \in G_0} t_Q, \quad w = \sum_{Q \in G_0} w_Q, \quad r(x) = x + \sum_{Q \in G_0} \left(\frac{t_Q}{x - x_Q} + \frac{u_Q}{(x - x_Q)^2} \right),$$

where G_0 is the set of nonzero points in G. The rational map $\phi(x, y) = (r(x), r'(x)y)$ is a separable isogeny from E to E' with equation $y^2 = x^3 + (A - 5t)x + B - 7w$ with kernel G.

We only considered the cases in which G has order 2 or an odd number. This covers all cases in which the degree is prime and from them we can build every isogeny, thanks to Theorem 1.2.24.

1.3 Dual Isogenies

In this section, we are going to define the concept of *dual isogeny* and deduce the structure of the torsion subgroups of elliptic curves from their properties.

Theorem 1.3.1. *Let $\phi \colon E \longrightarrow E'$ be a nonconstant isogeny of degree m. There exists a unique isogeny*

$$\widehat{\phi} \colon E' \longrightarrow E$$

satisfying $\widehat{\phi} \circ \phi = [m]$ and it is called dual isogeny to ϕ.

Proof. For the full proof, see [33, Theorem III.6.1]. Here, we are going to show only uniqueness. If $\widehat{\phi}$ and $\widehat{\phi}'$ are two isogenies satisfying the required property, then

$$(\widehat{\phi} - \widehat{\phi}') \circ \phi = [m] - [m] = [0].$$

Since ϕ is nonzero, $\widehat{\phi} - \widehat{\phi}'$ must be zero, so $\widehat{\phi} = \widehat{\phi}'$. □

Example 1.3.2. Let E and F be two elliptic curves and $\phi \colon E \longrightarrow F$ an isogeny between them, as in Example 1.2.9. The dual isogeny of ϕ is

$$\widehat{\phi} \colon F \longrightarrow E \quad \widehat{\phi}(x, y) = \left(\frac{5x^2 - 1}{x}, \frac{-7x^2y - 9y}{x^2} \right).$$

Since ϕ has degree 2, we have $\widehat{\phi} \circ \phi = [2]$.

Remark 1.3.3. Notice that also the constant isogeny has a "dual", but it is not unique. Indeed, every isogeny is the dual isogeny of $[0]$. It is customary to not include the constant isogeny in the definition in order to have both uniqueness and the following results.

In the following proposition, we summarize all the main properties of dual isogenies.

Proposition 1.3.4. *Let $\phi \colon E \longrightarrow E'$ be a nonconstant isogeny of degree m.*

(1) *We have $\widehat{\phi} \circ \phi = [m]$ on E and $\phi \circ \widehat{\phi} = [m]$ on E'.*

(2) *Let $\lambda \colon E' \longrightarrow E''$ be another isogeny. Then $\widehat{\lambda \circ \phi} = \widehat{\phi} \circ \widehat{\lambda}$.*

(3) *Let $\psi \colon E \longrightarrow E'$ be another isogeny. Then $\widehat{\phi + \psi} = \widehat{\phi} + \widehat{\psi}$.*

(4) *For all $m \in \mathbb{Z}$ we have $\widehat{[m]} = [m]$ and deg$[m] = m^2$.*

(5) *We have degϕ = deg$\widehat{\phi}$ and $\widehat{\widehat{\phi}} = \phi$.*

Proof. See [33, Theorem III.6.2]. □

These properties make the restriction of the dualization map

$$\widehat{\cdot} \colon \mathrm{End}(E) \longrightarrow \mathrm{End}(E)$$

an *anti-involution.*

Definition 1.3.5. Let E be an elliptic curve, $m \in \mathbb{Z}$. We define the *m-torsion subgroup* of E to be

$$E[m] = \{P \in E \colon [m]P = O\}.$$

Remark 1.3.6. Recall that every element in a finite group has finite order. We call them *torsion elements.* Let E be an elliptic curve defined over a finite field K of positive characteristic p. Then every point in E is a torsion point. Indeed, let $P \in E$. As

$$\overline{\mathbb{F}}_p = \bigcup_{n \in \mathbb{N}} \mathbb{F}_{p^n},$$

there exists $m \in \mathbb{N}$ such that P is $\mathbb{F}_{p^m}$-rational. Hence, P is contained in a finite subgroup of $E/\overline{\mathbb{F}_p}$, namely $E(\mathbb{F}_{p^m})$, and so we can conclude.

Corollary 1.3.7. *Let E be an elliptic curve and $m \in \mathbb{Z}$ with $m \neq 0$. If $m \neq 0$ in K, i.e. if either char$K = 0$ or $p =$ char$K > 0$ and p does not divide m, then*

$$E[m] = \frac{\mathbb{Z}}{m\mathbb{Z}} \times \frac{\mathbb{Z}}{m\mathbb{Z}}.$$

If $p =$ char$K > 0$, then one of the following is true:

(1) $E[p^i] = \{O\}$ *for all* $i = 1, 2, 3....$

(1) $E[p^i] = \dfrac{\mathbb{Z}}{p^i\mathbb{Z}}$ *for all* $i = 1, 2, 3....$

Proof. The assumption on m in the first claim and the fact that $\deg[m] = m^2$ imply that $[m]$ is a nonconstant separable isogeny. Hence, we have

$$\#E[m] = \# \ker[m] = \deg[m] = m^2.$$

Moreover, for each $d \in \mathbb{Z}$ dividing m, we have

$$\#E[d] = d^2$$

and these are subgroups of $E[m]$. Writing $E[m]$ as a product of cyclic groups, one can check that the only possibility is

$$E[m] = \frac{\mathbb{Z}}{m\mathbb{Z}} \times \frac{\mathbb{Z}}{m\mathbb{Z}}.$$

Let now π_p be the p-Frobenius morphism. Then we have

$$\#E[p^i] = \deg_s[p^i] = (\deg_s[p])^i = (\deg_s(\widehat{\phi} \circ \phi))^i = (\deg_s \widehat{\phi})^i,$$

since the Frobenius morphisms are purely inseparable. Moreover, we have

$$\deg \widehat{\phi} = \deg \phi = p,$$

so there are only two cases. If $\widehat{\phi}$ is inseparable, then $\deg_s \widehat{\phi} = 1$, so

$$\#E[p^i] = 1 \quad \text{for all} \quad i.$$

Otherwise, $\widehat{\phi}$ is separable, so $\deg_s \widehat{\phi} = p$ and

$$E[p^i] = \frac{\mathbb{Z}}{p^i\mathbb{Z}} \quad \text{for all} \quad i.$$

□

Definition 1.3.8. Let E be an elliptic curve defined over a field of characteristic $p > 3$. If $E[p] = \mathbb{Z}/p\mathbb{Z}$, then E is said to be *ordinary*. If $E[p] = \{O\}$, then E is said to be *supersingular*.

Example 1.3.9. The curve $E/\mathbb{F}_{19}$ with short Weierstrass equation $y^2 = x^3 + x$ is supersingular. The curve $F/\mathbb{F}_{19}$ with short Weierstrass equation $y^2 = x^3 + 12x + 12$ is ordinary. At this point, it is not trivial to understand whether an elliptic curve is supersingular or ordinary. However, we will see that the number of $\mathbb{F}_p$-rational points of an $\mathbb{F}_p$-rational supersingular elliptic curve defined over a finite field is exactly $p + 1$. Thus, in the above cases one can just check their cardinality in SageMath. The curve E has 20 points $\mathbb{F}_p$-rational points, while F has 15.

Lemma 1.3.10. *Any elliptic curve endomorphism $\phi \in \mathrm{End}(E)$ satisfies*

$$\phi + \widehat{\phi} = [1] + [\deg \phi] - [\deg([1] - \phi)].$$

Proof. We have

$$[\deg([1]-\phi)] = (\widehat{[1] - \phi})([1]-\phi) = (\widehat{[1]}-\widehat{\phi})([1]-\phi) = ([1]-\widehat{\phi})([1]-\phi) = [1]-(\phi+\widehat{\phi})+[\deg \phi],$$

which yields the desired relation. □

Notice that we can see both $\phi+\widehat{\phi}$ and $\widehat{\phi}\circ\phi$ as integers, even if they are endomorphisms.

Definition 1.3.11. Given E an elliptic curve and $\phi \in \mathrm{End}(E)$, we define its *trace* and its *norm* to be respectively

$$\mathrm{Tr}(\phi) = \phi + \widehat{\phi}, \quad \mathrm{N}(\phi) = \widehat{\phi} \circ \phi = [\deg \phi].$$

We will consider both the trace and the norm to be integers, rather than endomorphisms.

Theorem 1.3.12. *Let $\phi \in \mathrm{End}(E)$, for some E elliptic curve. Then, both ϕ and $\widehat{\phi}$ are roots of the polynomial*

$$x^2 - \mathrm{Tr}(\phi)x + \mathrm{N}(\phi) = 0.$$

Proof. One can directly check that we have

$$\phi^2 - (\phi + \widehat{\phi})\phi + \widehat{\phi}\phi = 0,$$

and the same argument applies for $\widehat{\phi}$. □

1.4 Hasse's Theorem

Let $E/\mathbb{F}_q$ be an elliptic curve over the finite field with $q = p^r$ elements, with p a prime natural number. In this section, our goal is to establish an upper bound for the number of $\mathbb{F}_q$ rational points of E. A first trivial upper bound can be found looking at the number of solution of the equation

$$y^2 = x^3 + Ax + B, \quad \text{with} \quad (x, y) \in \mathbb{F}_q^2.$$

Each value of x yields at most two values for y and we have to take into count also the point at infinity. Hence, the trivial upper bound is

$$\#E(\mathbb{F}_q) \leq 2q + 1.$$

In order to sharpen this estimate, we first introduce a lemma, which is similar to Cauchy-Schwarz inequality.

Lemma 1.4.1. *Let G be an abelian group and let $d\colon G \longrightarrow \mathbb{Z}$ be a positive definite quadratic form. Then*

$$|d(g-h) - d(h) - d(g)| \leq 2\sqrt{d(g)d(h)},$$

for all $g, h \in G$.

Proof. See [33, Lemma V.1.2]. □

Our definition of degree of an isogeny induces a degree map

$$\deg\colon \operatorname{End}(E) \longrightarrow \mathbb{Z}$$

and one can verify easily that it is a positive definite quadratic form on the abelian group $\operatorname{End}(E)$. The following result provides a much more precise upper bound. It was conjectured by E. Artin and then proved by Hasse.

Theorem 1.4.2 (Hasse)**.** *Let $E/\mathbb{F}_q$ be an elliptic curve over the finite field with $q = p^r$ elements, with p a prime natural number. Then we have*

$$\#E(\mathbb{F}_q) = q + 1 - t,$$

where t is the trace of the Frobenius endomorphism and $|t| \leq 2\sqrt{q}$.

Proof. The field $\mathbb{F}_q$ is the splitting field of $x^q - x$ over $\mathbb{F}_p$ and thus

$$\mathbb{F}_q = \{x \in \overline{\mathbb{F}_q}\colon x^q - x = 0\}$$

is precisely the subfield of $\overline{\mathbb{F}}_q$ fixed by the Frobenius automorphism $x \longmapsto x^q$. Hence, denoting by π the Frobenius endomorphism of E, we have that

$$E(\mathbb{F}_q) = \{P \in E\colon \pi(P) = P\} = \ker(\pi - [1]).$$

Recall that the Frobenius endomorphism is purely inseparable by Theorem 1.2.17 and $-[1]$ is separable, so by Proposition 1.2.21 we can conclude that $\pi - [1]$ is a separable map. Thus the cardinality of its kernel coincides with its degree, by Theorem 1.2.10. Therefore,

$$\#E(\mathbb{F}_q) = \#\ker(\pi - [1]) = \deg(\pi - [1]),$$

so we need to compute the degree of $\pi - [1]$. By Proposition 1.3.4, we have

$$[\deg(\pi - [1])] = (\widehat{\pi - [1]})(\pi - [1]) = \widehat{\pi}\pi + [1] - (\pi + \widehat{\pi}) = [q] + [1] - [t],$$

where t is the trace of the Frobenius endomorphism and q is its degree, or equivalently its norm. Thus, we have

$$\#E(\mathbb{F}_q) = q + 1 - t.$$

It remains to prove that $|t| \leq 2\sqrt{q}$. Notice that reordering the previous chain of equalities, we get

$$t = q + 1 - \#E(\mathbb{F}_q).$$

Moreover, we can write all the terms on the right hand side using the degree function and obtain

$$|t| = |\#E(\mathbb{F}_q) - q - 1| = |\deg(\pi - [1]) - \deg(\pi) - \deg([1])|.$$

Applying Lemma 1.4.1, we get

$$|t| \leq 2\sqrt{\deg(\pi)\deg([1])} = 2\sqrt{q},$$

as we wanted to prove. □

Remark 1.4.3. Hasse's upper bound is sharp, in the sense that it is the the best possible one. Indeed, when q is prime, there are elliptic curves over $\mathbb{F}_q$ with cardinalities matching every integer value in the *Hasse's interval*

$$\mathcal{H}(q) = [q + 1 - 2\sqrt{q}, q + 1 + 2\sqrt{q}].$$

For example, see [30, Theorem 4.2]

The next proposition shows that if we find elliptic curves for every value of the left half of the Hasse's interval (or equivalently the right half), then we automatically have curves for each value of the other half.

Proposition 1.4.4. *Let $\mathbb{F}_q$ be the field with q elements. Suppose that $u \in \mathbb{F}_q$ is not a square. Let $E/\mathbb{F}_q$ be the elliptic curve defined by the short Weierstrass equation*

$$y^2 = x^3 + Ax + B$$

and let E' be its quadratic twist by u. Then the traces of the Frobenius endomorphisms of E and E' are opposite integers.

Proof. Let $\chi\colon \mathbb{F}_q \longrightarrow \{-1, 0, 1\}$ be the function such that

$$\chi(x) = \begin{cases} -1 & \text{if } x \text{ is not a square in } \mathbb{F}_q, \\ 0 & \text{if } x \text{ is zero in } \mathbb{F}_q, \\ 1 & \text{if } x \text{ is a square in } \mathbb{F}_q. \end{cases}$$

Then the number of $\mathbb{F}_q$-rational points of E is

$$\#E(\mathbb{F}_q) = q + 1 + \sum_{x \in \mathbb{F}_q} \chi(x^3 + Ax + B),$$

therefore the trace of the Frobenius of E is $t = -\sum_{x \in \mathbb{F}_q} \chi(x)$. Recall that an equation for E' is

$$uy^2 = x^3 + Ax + B,$$

from which we can see that the point (x, y) is in E' if and only if $x^3 + Ax + B$ is not a square in $\mathbb{F}_q$. Hence, the number of $\mathbb{F}_q$-rational points of E' is

$$\#E(\mathbb{F}_q) = q + 1 - \sum_{x \in \mathbb{F}_q} \chi(x^3 + Ax + B),$$

from which we can conclude. □

The function χ of the above proof is usually denoted as $(\frac{x}{\mathbb{F}_q})$ and, if q is a prime, it coincides with the *Legendre symbol.* Indeed, the classical definition for the Legendre symbol is the following. If p is an odd prime and a is any integer, then

$$\left(\frac{a}{p}\right) = \begin{cases} -1 & \text{if } (a,p)=1 \text{ and } a \text{ is not a square modulo } p, \\ 0 & \text{if } p \text{ divides } a, \\ 1 & \text{if } (a,p)=1 \text{ and } a \text{ is a square modulo } p. \end{cases}$$

If $p = 2$, we define only for $a \equiv 0, 1$ modulo 4

$$\left(\frac{a}{p}\right) = \begin{cases} -1 & \text{if } a \equiv 5 \text{ modulo } 8, \\ 0 & \text{if } 4 \text{ divides } a, \\ 1 & \text{if } a \equiv 1 \text{ modulo } 8. \end{cases}$$

We will apply the Legendre symbol in the case a is the discriminant of a quadratic number field and those are either congruent to 0 or 1 modulo 4, hence we are not leaving out any useful possibility. We recall the following standard result in number theory: a prime number p in a quadratic number field of discriminant d is split, ramified or inert if and only if $(\frac{d}{p})$ is 1, 0 or -1, respectively.

To conclude this section, we present a classical result on finite fields elliptic curves, due to Tate.

Theorem 1.4.5 (Tate)**.** *Two elliptic curves E, E' defined over a finite field $\mathbb{F}_q$ are isogenous over $\mathbb{F}_q$ if and only if their traces of Frobenius endomorphisms are equal.*

Proof. See [34] for a proof valid for general abelian varieties. □

Remark 1.4.6. By Theorem 1.4.2, given E an elliptic curve over $\mathbb{F}_q$, we know that the trace of the Frobenius endomorphism depends only on the number of $\mathbb{F}_q$-rational points. Hence, we can rephrase Tate Theorem in an equivalent way.

Theorem 1.4.7 (Tate)**.** *Two elliptic curves E, E' defined over a finite field $\mathbb{F}_q$ are isogenous over $\mathbb{F}_q$ if and only if $\#E(\mathbb{F}_q) = \#E'(\mathbb{F}_q)$.*

1.5 Endomorphism Algebras

In this section we aim to get a better understanding of the structure of the endomorphism ring of an elliptic curve. Let E/K be an elliptic curve. The ring $\mathrm{End}(E)$ needs not to be commutative, but we know that its center contains all multiplication by m maps, for $m \in \mathbb{Z}$. Moreover, it forms a subring of $\mathrm{End}(E)$. In the following, we will identify this subring with $\mathbb{Z}$ and we may write m rather than $[m]$ to ease the notation, when our intentions are clear from the context.

We recall one possible definition for the tensor product, which is practical and does not involve universal properties.

Definition 1.5.1. Let R be a commutative ring and A, B two R-modules. Their *tensor product* over R, denoted by $A \otimes_R B$, is the R-module generated by the formal symbols $\alpha \otimes \beta$, where $\alpha \in A$ and $\beta \in B$, with the relations

$$(\alpha_1+\alpha_2)\otimes\beta = \alpha_1\otimes\beta+\alpha_2\otimes\beta, \quad \alpha\otimes(\beta_1+\beta_2) = \alpha\otimes\beta_1+\alpha\otimes\beta_2, \quad r\alpha\otimes\beta = \alpha\otimes r\beta = r(\alpha\otimes\beta),$$

for all $\alpha_1, \alpha_2 \in A$ and $\beta_1, \beta_2 \in B$.

Definition 1.5.2. Let E/K be an elliptic curve. Its *endomorphism algebra* is

$$\mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}$$

and we denote it as $\mathrm{End}^A(E)$.

In order to give a uniform way to denote elements of the endomorphism algebra, we use the following lemma.

Lemma 1.5.3. *Let R be an integral domain with fraction field B and let A be an R-algebra. Every element of $A \otimes_R B$ can be written in the form $\alpha \otimes \beta$ with $\alpha \in A$ and $\beta \in B$.*

Proof. It is sufficient to show that, for every $\alpha_1\alpha_2 \in A$ and every $\beta_1, \beta_2 \in B$, the element $\alpha_1 \otimes \beta_1 + \alpha_2 \otimes \beta_2$ can be written as $\alpha_3 \otimes \beta_3$, for some $\alpha_3 \in A$ and $\beta_3 \in B$. Let $\beta_1 = r_1/s_1$ and $\beta_2 = r_2/s_2$, for some $r_1, r_2, s_1, s_2 \in R$. Then, we have

$$\begin{aligned}
\alpha_1 \otimes \beta_1 + \alpha_2 \otimes \beta_2 =& \alpha_1 \otimes \frac{r_1}{s_1} + \alpha_2 \otimes \frac{r_2}{s_2} \\
=& \alpha_1 \otimes \frac{r_1 s_2}{s_1 s_2} + \alpha_2 \otimes \frac{r_2 s_1}{s_1 s_2} \\
=& r_1 s_2 \alpha_1 \otimes \frac{1}{s_1 s_2} + r_2 s_1 \alpha_2 \otimes \frac{1}{s_1 s_2} \\
=& (r_1 s_2 \alpha_1 + r_2 s_1 \alpha_2) \otimes \frac{1}{s_1 s_2}.
\end{aligned}$$

Hence, we can take $\alpha_3 = r_1 s_2 \alpha_1 + r_2 s_1 \alpha_2$ and $\beta_3 = 1/(s_1 s_2)$. □

The fact that $\mathrm{End}(E)$ has a multiplication compatible with its $\mathbb{Z}$-module structure makes it a $\mathbb{Z}$-algebra. Thus, the lemma implies that every element of $\mathrm{End}^A(E)$ can be written as $\phi \otimes r$ for some $\phi \in \mathrm{End}(E)$ and $r \in \mathbb{Q}$. In order to ease the notation we may write $r\phi$. Notice that, for $m \in \mathbb{Z}$, the difference between $r\phi$ and $m\phi$ is that the former is not necessarily an endomorphism, but it is a formal element of $\mathrm{End}^A(E)$.

Next, we want to extend the involution $\phi \longmapsto \widehat{\phi}$, the trace and the norm maps to $\mathrm{End}^A(E)$.

Definition 1.5.4. Let ϕ be an endomorphism of E. We define

$$\widehat{r\phi} = r\widehat{\phi} \quad \text{for all} \quad r \in \mathbb{Q}.$$

Let $\alpha \in \mathrm{End}^A(E)$. We define the norm of α and the trace of α to be respectively

$$\mathrm{N}(\alpha) = \alpha\widehat{\alpha} \quad \text{and} \quad \mathrm{Tr}(\alpha) = \alpha + \widehat{\alpha}.$$

Under this definition, the dualization map preserves all its nice properties and therefore it is still an anti-involution of $\operatorname{End}^A(E)$, which is customarily called *Rosati involution.* Moreover, trace and norm maps have the following properties.

Proposition 1.5.5. *Let E/K be an elliptic curve, $\alpha, \beta \in \operatorname{End}^A(E)$ and $r \in \mathbb{Q}$.*

(1) $\mathrm{N}(\alpha) \in \mathbb{Q}$ *and* $\mathrm{N}(\alpha) \geq 0$*. Moreover,* $\mathrm{N}(\alpha) = 0$ *if and only if* $\alpha = 0$*.*

(2) $\mathrm{N}(\alpha) = \mathrm{N}(\widehat{\alpha})$ *and* $\mathrm{N}(\alpha\beta) = \mathrm{N}(\alpha)\,\mathrm{N}(\beta)$*.*

(3) *If α is nonzero, then it admits a multiplicative inverse.*

(4) $\mathrm{Tr}(\alpha) = \mathrm{Tr}(\widehat{\alpha}) \in \mathbb{Q}$*. Moreover* $\mathrm{Tr}(\alpha + \beta) = \mathrm{Tr}(\alpha) + \mathrm{Tr}(\beta)$ *and* $\mathrm{Tr}(r\alpha) = r\,\mathrm{Tr}(\alpha)$*.*

(5) *α and $\widehat{\alpha}$ are the roots of the polynomial $x^2 - \mathrm{Tr}(\alpha)x + \mathrm{N}(\alpha) \in \mathbb{Q}[x]$.*

(6) *Assume α is nonzero. If $\mathrm{Tr}(\alpha) = 0$, then $\alpha^2 = -\mathrm{N}(\alpha) < 0$.*

(7) *α is fixed by the Rosati involution if and only if $\alpha \in \mathbb{Q}$.*

Proof. All these properties can be proved by trivial computations. Here, we only provide proofs for the third and the seventh. Let us assume that α is a nonzero element of the endomorphism algebra. Then, using the first property, we have that $\mathrm{N}(\alpha) \neq 0$. Let $\gamma = \widehat{\alpha}/\mathrm{N}(\alpha)$. We have $\alpha\gamma = 1$, so α is invertible. By definition, we have $\widehat{r} = r$ for each $r \in \mathbb{Q}$. Assume that $\alpha = \widehat{\alpha}$. Then, exploiting the sixth property, we have that the discriminant $\mathrm{Tr}(\alpha)^2 - 4\,\mathrm{N}(\alpha)$ of the polynomial is zero, in which case $\alpha = \frac{1}{2}\,\mathrm{Tr}(\alpha) \in \mathbb{Q}$. □

Remark 1.5.6. The third statement of the previous proposition implies that $\operatorname{End}^A(E)$ is a division ring. This means that it is a field if and only if the multiplication is commutative.

In order to give a complete classification of the endomorphism algebras of elliptic curve, we need to briefly recall the notions of quadratic imaginary number field and quaternion algebra.

An imaginary quadratic number field is a field extension $\mathbb{Q}(\sqrt{m})$ of the rational numbers of degree two, with $m < 0$. Its ring of integers is the set of all complex numbers that are roots of monic polynomials with coefficients in $\mathbb{Z}$.

A quaternion algebra over a field K is a K-algebra $\mathbb{Q}(i, j)$ that has a basis of the shape $\{1, i, j, k\}$ such that

$$i^2, j^2 \in K^\times \quad \text{and} \quad ij = k = -ji.$$

We are now ready to introduce the aforementioned classification.

Theorem 1.5.7. *Let E/K be an elliptic curve. Then $\operatorname{End}^A(E)$ is isomorphic to one of the following:*

(1) *the field of rational numbers;*

(2) *an imaginary quadratic number field;*

(3) *a quaternion algebra over the rational numbers $\mathbb{Q}(i, j)$, with $i^2, j^2 < 0$.*

Proof. First, observe that we always have $\mathbb{Q} \subseteq \mathrm{End}^A(E)$. If equality holds, then we fall into the first option given by the theorem. Now, suppose that the inclusion is strict. Let $i \in \mathrm{End}^A(E)$ such that $i \notin \mathbb{Q}$. We may assume without loss of generality that $\mathrm{Tr}(i) = 0$. Indeed we can replace i with the element $i - \frac{1}{2}\mathrm{Tr}(i)$. By the sixth property of Lemma 1.5.5, we have that $i^2 < 0$. This yields the inclusion $\mathbb{Q}(i) \subseteq \mathrm{End}^A(E)$. If equality holds, we fall into the second option given by the theorem. Otherwise, assume the inclusion is strict. Let j be an element of $\mathrm{End}^A(E)$ not lying in $\mathbb{Q}(i)$. Again, we can assume without loss of generality that $\mathrm{Tr}(j) = 0$, for the same exact reason as above. This yields $j^2 < 0$. Furthermore, we may replace j with

$$j - \frac{\mathrm{Tr}(ij)}{2i^2} i$$

and so we can assume $\mathrm{Tr}(ij) = 0$. Indeed, one could directly compute the trace and check that this holds. Moreover, this alteration does not change the trace of j, since $\mathrm{Tr}(i) = 0$. To summarize, we have $\mathrm{Tr}(i) = \mathrm{Tr}(j) = \mathrm{Tr}(ij) = 0$ and this implies $i = -\widehat{i}, j = -\widehat{j}$ and $ij = -\widehat{ij} = -\widehat{\widehat{i}\widehat{j}}$. Combining these equalities yields $ij = -ji$. Let us define $k = ij$. It is clear that the set $\{1, i, j, k\}$ spans the whole $\mathbb{Q}(i, j)$ as a $\mathbb{Q}$-vector space, but, if we want to prove it is a basis, we need to prove that its elements are linearly independent. By construction, 1, i and j are linearly independent. Suppose by contradiction that there exist $a, b, c \in \mathbb{Q}$ such that

$$k = a + bi + cj.$$

Observe we need to have $c \neq 0$, since ij does not lie in $\mathbb{Q}(i)$. Taking squares of both sides yields

$$(ij)^2 = (a^2 + b^2 i^2 + c^2 j^2) + 2a(bi + cj) + bc(ij + ji).$$

The left hand side and the first term of the right hand side lie in $\mathbb{Q}$, since $\mathrm{Tr}(i) = \mathrm{Tr}(j) = \mathrm{Tr}(ij) = 0$. The last term of the right hand side is zero, since $ij = -ji$. Hence, the term $x = bi + cj$, for some $x \in \mathbb{Q}$. This yields $j = (x - bi)/c \in \mathbb{Q}(i)$, which is a contradiction. If $\mathbb{Q}(i, j) = \mathrm{End}^A(E)$, we fall into the third option given by the theorem. Let us suppose by contradiction that the inclusion is strict. Let h be an element of $\mathrm{End}^A(E)$ not lying in $\mathbb{Q}(i, j)$. Again, we may assume without loss of generality that $\mathrm{Tr}(h) = \mathrm{Tr}(ih) = 0$, which implies $ih = -hi$. Then, we have $ijh = -jih = jhi$, so i commutes with jh. Lemma 1.5.8 implies $jh \in \mathbb{Q}(i)$, which implies $h \in \mathbb{Q}(i, j)$, contrary to our hypothesis. □

Lemma 1.5.8. *If $x, y \in \mathrm{End}^A(E)$ commute and $x \notin \mathbb{Q}$, then $y \in \mathbb{Q}(x)$.*

Proof. Let x' and y' be the elements obtained combining $\mathbb{Q}$-linearly x and y such that $\mathrm{Tr}(x') = \mathrm{Tr}(y') = \mathrm{Tr}(x'y') = 0$, so that $x'y' = -y'x'$. This can be done in the following way

$$x' = x - t, \quad y' = y - s - tx,$$

for suitable $r, s, t \in \mathbb{Q}$ as in the previous proof. If x and y commute, then so do all their linear combinations. We have $2x'y' = 0$, which implies $x' = 0$ or $y' = 0$, as $\mathrm{End}^A(E)$ is a division ring. Since $x' \notin \mathbb{Q}$, we must have $y' = 0$. This implies that $y = s + tx \in \mathbb{Q}(x)$, as we wanted. □

If we want to understand the structure of the endomorphism ring, we need to introduce the following definition.

Definition 1.5.9. Let A be a $\mathbb{Q}$-algebra of finite dimension r as a $\mathbb{Q}$-vector space. An *order* $\mathcal{O}$ in A is a subring of A such that $A = \mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q}$. Equivalently, an order is a subring of A that is a rank r free $\mathbb{Z}$-module.

By definition of $\operatorname{End}^A(E)$, the endomorphism ring is an order in the endomorphism algebra. Hence, we have the following corollary.

Corollary 1.5.10. *Let E/K be an elliptic curve. One of the following holds.*

1. $\operatorname{End}(E) \simeq \mathbb{Z}$.

2. $\operatorname{End}(E)$ *is an order in a quadratic imaginary number field.*

3. $\operatorname{End}(E)$ *is an order in a quaternion algebra.*

Proof. The claim follows directly from the classification Theorem 1.5.7, after noticing that every order in $\mathbb{Q}$ is isomorphic to $\mathbb{Z}$. □

Definition 1.5.11. An elliptic curve whose endomorphism ring $\mathcal{O}$ strictly contains $\mathbb{Z}$ is said to have *complex multiplication* by $\mathcal{O}$.

One can prove that the ring of integers of any number field is its unique maximal order. Indeed, this follows from the fact that any order in a number field L is a subring that contains a $\mathbb{Q}$-basis of L.

Furthermore, one can give a complete characterization of orders in quadratic imaginary number fields.

Lemma 1.5.12. *Let L be an imaginary quadratic number field with ring of integers $\mathcal{O}_L$ and let $\mathcal{O}$ be an order in L. Then there exists a unique positive integer f such that*

$$\mathcal{O} = \mathbb{Z} + f\mathcal{O}_L.$$

The integer f is called the conductor of the order $\mathcal{O}$ and it is equal to the index $[\mathcal{O}_L \colon \mathcal{O}]$.

Proof. See [10, Lemma 7.2]. □

In general, it is nontrivial to determine the structure of the endomorphism ring of an elliptic curve given in short Weierstrass equation. On the other hand, the structure of the *automorphism group* is much simpler and we conclude this section with a theorem on its structure.

Theorem 1.5.13. *Let E/K be an elliptic curve with j-invariant j. If j is different from* 0 *and* 1728, *then* $\operatorname{Aut}(E)$ *is isomorphic to* $\mathbb{Z}/2\mathbb{Z}$. *If $j = 0$ or $j = 1728$, then* $\operatorname{Aut}(E)$ *has cardinality* 6 *and* 4, *respectively.*

Proof. Let $y^2 = x^3 + Ax + B$ be a short Weierstrass equation for E. Then, by Theorem 1.1.7, every automorphism of E has the form

$$x = u^2 x', \quad y = u^3 y',$$

for some $u \in \overline{K}$ and this gives an automorphism if and only if

$$A = u^{-4} A, \quad \text{and} \quad B = u^{-6} B.$$

If $AB \neq 0$, which means that $j(E) \neq 0, 1728$, then the only possibilities are $u = \pm 1$.

If $A = 0$, then $j(E) = 0$ and we must have $u^6 = 1$. Similarly, if $B = 0$, then $j(E) = 1728$ and $u^4 = 1$. Hence, $\mathrm{Aut}(E)$ is cyclic of order 2, 4 or 6, depending on the j-invariant of E. □

1.6 Supersingular and Ordinary Elliptic Curves

Recall that in Definition 1.3.8 we distinguished ordinary and supersingular elliptic curves over finite fields of characteristic p according to the structure of their p-torsion subgroups. We begin this section by proving that being supersingular is an isogeny-invariant property.

Proposition 1.6.1. *Let $\phi\colon E_1 \longrightarrow E_2$ be an isogeny over a field of positive characteristic p. Then E_1 is supersingular (ordinary) if and only if E_2 is supersingular (ordinary).*

Proof. Let p_1 and p_2 denote the multiplication by p maps on E_1 and E_2, respectively. We have

$$p_2 \circ \phi = \phi + \ldots + \phi = \phi \circ p_1,$$

and so

$$\deg_s(p_2) = \deg_s(p_1).$$

E_1 is supersingular if and only if $E[p] = \{O\}$, i.e. if and only if the map p_1 is purely inseparable. The theorem follows from the equality of separable degrees. □

The following theorem provides a useful characterization of supersingular curves.

Theorem 1.6.2. *An elliptic curve $E/\mathbb{F}_q$ is supersingular if and only if $\mathrm{Tr}(\pi) \equiv 0 \bmod p$, where π is the Frobenius endomorphism of E and p is the characteristic of $\mathbb{F}_q$.*

Proof. Let $q = p^n$, for some n natural number. Recall that all p^r-Frobenius morphisms are purely inseparable, by Lemma 1.2.17. Thus, we have

$$\deg_s(p) = \deg_s(\pi_p) \deg_s(\widehat{\pi_p}) = \deg_s(\widehat{\pi_p}).$$

E is supersingular if and only if $E[p] = \ker[p] = \deg_s(p) = \deg_s(\widehat{\pi_p})$ is trivial. Hence, E is supersingular if and only if $\widehat{\pi_p}$ is purely inseparable.

The isogeny $\hat{\pi} = \widehat{\pi_p^n} = \widehat{\pi_p}^n$ is also purely inseparable. Therefore, $\mathrm{Tr}(\pi) = \pi + \hat{\pi}$ is the sum of purely inseparable isogenies and so it is itself a purely inseparable isogeny, as we proved in Proposition 1.2.21. Thus, $\mathrm{Tr}(\pi) \equiv 0 \bmod p$ because $[\mathrm{Tr}(\pi)]$ is inseparable if and only if p divides $[\mathrm{Tr}(\pi)]$.

Conversely, if $\mathrm{Tr}(\pi) \equiv 0 \bmod p$, then $[\mathrm{Tr}(\pi)]$ is purely inseparable and so does $\widehat{\pi} = \mathrm{Tr}(\pi) - \pi$. This implies that $\widehat{\pi_p}^n$ and $\widehat{\pi_p}$ are purely inseparable. Thus, the theorem holds. □

This theorem yields a corollary which is extremely important in the applications. It is the main feature that makes the cryptosystem CSIDH really fast.

Corollary 1.6.3. *Let $E/\mathbb{F}_p$ be an elliptic curve over a field of prime order $p > 3$ and let π be its Frobenius endomorphism. Then E is supersingular if and only if $\operatorname{Tr}(\pi) = 0$ or, equivalently, if and only if $\#E(\mathbb{F}_p) = p + 1$.*

Proof. By the Hasse's theorem, $|\operatorname{Tr}(\pi)| \leq 2\sqrt{p}$, and if $p > 3$ we have $2\sqrt{p} < p$. Hence the only possibility is that the trace of the Frobenius endomorphism is zero. □

The property of being supersingular refletcs also on the j-invariant, reducing its possible values.

Theorem 1.6.4. *Let E/K be a supersingular elliptic curve defined over a finite field of positive characteristic p. Then $j(E)$ lies in $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$.*

Proof. Assume E is in short Weierstrass form

$$y^2 = x^3 + Ax + B.$$

For each q power of p, let $E^{(q)}$ be the elliptic curve whose short Weierstrass equation is

$$y^2 = x^3 + A^q x + B^q,$$

and π_p as in Definition 1.2.16. Since E is supersingular, the isogeny $\widehat{\pi_p}\colon E^{(p)} \longrightarrow E$ is purely inseparable of degree p. By Lemma 1.2.19, we can write $\widehat{\pi_p} = \psi \circ \pi_p$, for some ψ separable isogeny which, in this case, must have degree 1. Thus, we have

$$p = \widehat{\pi_p} \circ \pi_p = \psi \circ \pi_p^2$$

and it follows that ψ is an isomorphism of $E^{(p^2)}$ to E. We have

$$j(E) = j(E^{(p^2)}) = j(E)^{p^2},$$

which implies that $j(E)$ is fixed by the automorphism $x \longmapsto x^{p^2}$, from which we can conclude. □

The above theorem should convince ourselves that the number of supersingular curves is really small compared to the number of ordinary curves. That is the reason why they are called supersingular, as a synonym of exceptional.

The next result characterizes all supersingular curves in characteristic $p > 3$.

Theorem 1.6.5. *Let $\mathbb{F}_q$ be a finite field with characteristic $p > 3$ and let E be an elliptic curve given by the equation*

$$y^2 = f(x),$$

for some $f \in \mathbb{F}_q[x]$. Then E is supersingular if and only if the coefficient of x^{p-1} in $f(x)^{(p-1)/2}$ is zero.

Proof. See [33, Theorem V.4.1]. □

We can count precisely the number of supersingular curves depending on the characteristic of the base field.

Corollary 1.6.6. *For $p \geq 5$ prime, the number (up to isomorphism) of supersingular elliptic curves defined over a finite field of characteristic p is*

$$\left\lfloor \frac{p}{12} \right\rfloor + \begin{cases} 0 & \text{if } p \equiv 1 \bmod 12, \\ 1 & \text{if } p \equiv 5, 7 \bmod 12, \\ 2 & \text{if } p \equiv 11 \bmod 12. \end{cases}$$

Proof. See [33, Theorem V.4.1]. □

Next, we give another characterization of supersingular and ordinary elliptic curves, perhaps the most important one.

Theorem 1.6.7. *If $E/\mathbb{F}_q$ is an ordinary elliptic curve, then $\mathrm{End}^A(E) = \mathbb{Q}(\pi)$ is an imaginary quadratic number field, where π is the Frobenius endomorphism. If $E/\mathbb{F}_q$ is a supersingular elliptic curve, then $\mathrm{End}^A(E)$ is a quaternion algebra.*

Proof. See [33, Theorem V.3.1.]. □

If $E/\mathbb{F}_q$ is an ordinary elliptic curve, then its Frobenius endomorphism is not an integer. Hence, the subring $\mathbb{Z}[\pi]$ is lattice of rank 2. Here, by lattice we mean a free $\mathbb{Z}$-module, and its rank is the cardinality of any set of linearly independent generators. Therefore, $\mathbb{Z}[\pi]$ is an order in $L = \mathrm{End}^A(E)$. However, the endomorphism ring of E does not need to be equal to $\mathbb{Z}[\pi]$. The fact that $\mathrm{End}(E)$ contains $\mathbb{Z}[\pi]$ and is contained in the maximal order $\mathcal{O}_L$ reduces its possible configurations to a finite number. We express this constraint in terms of conductors.

Proposition 1.6.8. *Let $E/\mathbb{F}_q$ be an ordinary elliptic curve, L its endomorpshism algebra, which is an imaginary quadratic number field. Let $f = [\mathcal{O}_L \colon \mathrm{End}(E)]$ and $f_\pi = [\mathcal{O}_L \colon \mathbb{Z}[\pi]]$ be the conductors respectively of $\mathrm{End}(E)$ and $\mathbb{Z}[\pi]$. Then f divides f_π.*

Proof. We must have

$$f_\pi = [\mathcal{O}_L \colon \mathbb{Z}[\pi]] = [\mathcal{O}_L \colon \mathrm{End}(E)][\mathrm{End}(E) \colon \mathbb{Z}[\pi]] = f[\mathrm{End}(E) \colon \mathbb{Z}[\pi]],$$

as we wanted. □

Corollary 1.6.9. *If $E/\mathbb{F}_q$ is an ordinary elliptic curve, then $\mathrm{End}^A(E) \simeq \mathbb{Q}(\sqrt{d})$, where $d = t^2 - 4q < 0$ is the discriminant of the order $\mathbb{Z}[\pi]$, with $t = \mathrm{Tr}(\pi)$.*

Proof. One can easily compute the discriminant of $\mathbb{Z}[\pi]$ using the properties of the Frobenius endomorphism:

$$d(\mathbb{Z}[\pi]) = \det \begin{pmatrix} 1 & \pi \\ 1 & \widehat{\pi} \end{pmatrix}^2 = (\widehat{\pi}-\pi)^2 = \widehat{\pi}^2 - 2\widehat{\pi}\pi + \pi^2 = t\widehat{\pi} - q - 2q + t\pi - q = t(\widehat{\pi}+\pi) - 4q = t^2 - 4q.$$

Moreover, by the previous theorem, $\mathrm{End}^A(E) = \mathbb{Q}(\pi)$. Since π is a root of the equation

$$x^2 - tx + q,$$

we have $\mathrm{End}^A(E) = \mathbb{Q}(\sqrt{d})$. Notice that by Hasse's theorem $d < 0$. □

Chapter 2

Complex Multiplication

In this chapter, we will go through complex multiplication theory. In general, it is the study of elliptic curves with complex multiplication, which are curves with endomorphism ring strictly bigger than $\mathbb{Z}$. In particular, we will focus on elliptic curves with complex multiplication by an order in an imaginary quadratic number field. We will show that there is a strong correspondence between complex lattices and complex elliptic curves. We will not go through all the details since this would require introducing complex analytic tools, which are somehow far from our main topic. The most important feature of this correspondence is that it also has an impact at the level of lattice inclusions and isogenies. Moreover, we are going to define the action of a group on some set of elliptic curves. This action will allow us to encode important information about isogenies.

2.1 Complex Elliptic Curves and Lattices

We start this section by recalling the definition of complex lattice.

Definition 2.1.1. A *complex lattice* Λ is a discrete additive subgroup of $\mathbb{C}$ that contains an $\mathbb{R}$-basis for $\mathbb{C}$. Let $\Lambda' \subseteq \Lambda$ be a subgroup such that the quotient group Λ'/Λ is cyclic. We say that Λ' is a *cyclic sublattice* of Λ.

In other words, a complex lattice is a discrete subset of $\mathbb{C}$ that is also a free $\mathbb{Z}$-module of rank 2. Hence, each complex lattice Λ can be identified with two complex numbers ω_1, ω_2. We have

$$\Lambda = \omega_1\mathbb{Z} + \omega_2\mathbb{Z}$$

and we will write $\Lambda = [\omega_1, \omega_2]$. The set $\{\omega_1, \omega_2\}$ is called a basis of Λ and it is not unique.

Since every complex lattice Λ is an additive subgroup of $\mathbb{C}$, we can consider the quotient group $\mathbb{C}/\Lambda$. This is also called *complex torus*, since one can show it is homeomorphic to a torus.

Definition 2.1.2. A *fundamental domain* for $\Lambda = [\omega_1, \omega_2]$ is any set of the form

$$\{\alpha + t_1\omega_1 + t_2\omega_2 \colon 0 \leq t_1, t_2 < 1\},$$

with $\alpha \in \mathbb{C}$ fixed.

Every point of $\mathbb{C}/\Lambda$ has a representative in any fundamental domain, and $\mathbb{C}/\Lambda$ can be identified with any fundamental domain.

Definition 2.1.3. Let $\Lambda \subseteq \mathbb{C}$ be a complex lattice. The *Weierstrass $\wp$-function* is defined by the series

$$\wp(z;\Lambda) = \frac{1}{z^2} + \sum_{\omega\in\Lambda,\omega\neq 0} \left(\frac{1}{(z-\omega)^2} + \frac{1}{\omega^2} \right).$$

The *Eisenstein series* of weight $2k$ is the series

$$G_{2k}(\Lambda) = \sum_{\omega\in\Lambda,\omega\neq 0} \omega^{-2k}.$$

If the complex lattice is clear from the context, we will omit the dependence from Λ and write simply $\wp(z)$ and G_{2k}.

The next theorem ensures that the series we have just introduced are well defined.

Theorem 2.1.4. *Let $\Lambda \subseteq \mathbb{C}$ be a complex lattice.*

(1) *The Eisenstein series G_{2k} is absolutely convergent for all $k > 1$.*

(2) *The series defining the Weierstrass $\wp$-function converges absolutely and uniformly on every compact subset of $\mathbb{C} \smallsetminus \Lambda$. The series defines a meromorphic function on $\mathbb{C}$ having a double pole with residue 0 at each point of Λ and no other poles.*

Proof. See [33, Theorem VI.3.1]. □

From the following theorem, we can begin foreseeing the relation between complex lattices and complex elliptic curves.

Theorem 2.1.5. *Let $\Lambda \subseteq \mathbb{C}$ be a complex lattice. For all $z \in \mathbb{C} \smallsetminus \Lambda$, the Weierstrass $\wp$-function and its derivative satisfy the relation*

$$\wp'(z)^2 = 4\wp(z)^3 - 60G_4\wp(z) - 140G_6.$$

Proof. See [33, Theorem VI.3.5]. □

We can observe that the above relation really resembles the short Weierstrass equation for elliptic curves, up to the coefficient of the cubic term. In order to ease the notation, it is customary to set

$$g_2 = g_2(\Lambda) = 60G_4(\Lambda) \quad \text{and} \quad g_3 = g_3(\Lambda) = 140G_6(\Lambda).$$

If we set $y = \wp'(z)$ and $x = \wp(z)$, the above differential relation corresponds to the curve

$$y^2 = 4x^3 - g_2x - g_3, \tag{2.1}$$

which can be easily put in short Weierstrass form with $g_2 = -4A$ and $g_3 = -4B$. Thus, every complex lattice Λ yields an equation for a potential complex elliptic curve, provided it is nonsingular. Suppose that the resulting curve is singular. Then, by definition, the partial derivatives of the equation

$$zy^2 = 4x^3 - g_2xz^2 - g_3z^3$$

vanish simultaneously at some point. This means that there is a projective solution to the system of equations

$$\begin{cases} 12x^2 - g_2z^2 = 0 \\ 2zy = 0 \\ y^2 + 2g_2xz + 3g_3z^2 = 0. \end{cases}$$

Without loss of generality, we can assume $z = 1$, since $z = 0$ yields no projective solution. The second equation implies $y = 0$, which forces $x = -3g_3/(2g_2)$. Plugging these relations into the first equation gives $g_2^3 - 27g_3^2 = 0$. Hence, as long as

$$\Delta(\Lambda) = g_2^3 - 27g_3^2$$

is nonzero, Equation (2.1) defines a complex elliptic curve.

Lemma 2.1.6. *Let $\Lambda \subseteq \mathbb{C}$ be a complex lattice. Then $\Delta(\Lambda) \neq 0$.*

Proof. See [33, Proposition VI.3.6]. □

From the previous discussion, it follows that we can define an elliptic curve associated to Λ and we will denote it by E_Λ. We have a natural map

$$\Phi\colon \mathbb{C}/\Lambda \longrightarrow E_\Lambda, \quad z \longmapsto \begin{cases} (\wp(z), \wp'(z)) & \text{if } z \notin \Lambda \\ 0 & \text{else.} \end{cases}$$

Theorem 2.1.7. *Let $\Lambda \subseteq \mathbb{C}$ be a complex lattice and E_Λ its associated elliptic curve. The map Φ is a group isomorphism.*

Proof. See [33, Proposition VI.3.6]. One can also prove that this map preserves the complex structure, i.e. it is an isomorphism of complex manifolds. □

Next, guided by the correspondence between complex elliptic curves and complex lattices, we give the following definition.

Definition 2.1.8. Let $\Lambda \subseteq \mathbb{C}$ be a complex lattice. Its *j-invariant* is defined as

$$j(\Lambda) = 1728\frac{g_2^3}{\Delta(\Lambda)}.$$

Notice that it is well defined, since every complex lattice has nonzero discriminant, by Lemma 2.1.6. As we have already noticed, the elliptic curve E_Λ is isomorphic to the elliptic curve $y^2 = x^3 + Ax + B$, where $g_2 = -4A$ and $g_3 = -4B$. Thus, we have

$$j(\Lambda) = 1728\frac{g_2^3}{g_2^3 - 27g_3^2} = 1728\frac{(-4A)^3}{(-4A)^3 - 27(-4B)^2} = 1728\frac{4A^3}{4A^3 + 27B^2} = j(E_\Lambda).$$

Definition 2.1.9. Lattices Λ_1 and Λ_2 are said to be homothetic if $\Lambda_2 = z\Lambda_1$, for some $z \in \mathbb{C}^*$.

Theorem 2.1.10. *Two complex lattices Λ_1 and Λ_2 are homothetic if and only if $j(\Lambda_1) = j(\Lambda_2)$.*

Proof. See [1, Theorem 1.16]. □

Example 2.1.11. Let us consider the complex lattices $\Lambda_1 = [1+i\frac{1}{5}, \frac{2}{5}+i\frac{7}{10}] = [\omega_1, \omega_2]$ and $\Lambda_2 = [\frac{4}{5}+i\frac{6}{5}, -\frac{3}{10}+i\frac{11}{10}] = [\tau_1, \tau_2]$, as in Figure 2.1. These are homothetic: the quotient of their former generators ω_1/τ_1 coincides with the quotient of their latter generator ω_2/τ_2. This also implies that their j-invariants are equal.

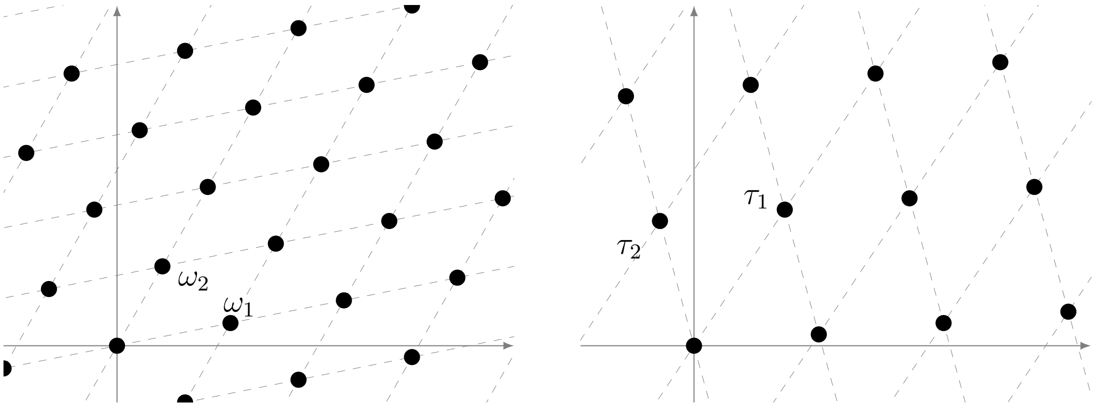

Figure 2.1: Two homothetic complex lattices.

Recall that two elliptic curves are isomorphic if and only if they have the same j-invariant. Hence, the previous theorem directly yields the following corollary.

Corollary 2.1.12. *Two complex lattices Λ_1 and Λ_2 are homothetic if and only if the elliptic curves E_{Λ_1} and E_{Λ_2} are isomorphic.*

Observe that homothety of lattices is an equivalent relation and, if we consider lattices up to homothety, each equivalence class has a representative of the shape $[1, \omega]$, for some $\omega \in \mathbb{C}$ with $\Im(\omega) > 0$. Indeed, for any complex lattice $\Lambda = [\omega_1, \omega_2]$, the complex lattice $\Lambda' = \omega_1^{-1}[\omega_1, \pm\omega_2] = [1, \pm\omega_1^{-1}\omega_2]$ is in the same equivalence class of Λ and in the desired form. Hence, we have a natural function

$$j \colon \mathbb{H} \longrightarrow \mathbb{C}, \quad \omega \longmapsto j([1, \omega]),$$

which associates to each point in the open upper half of the complex plane $\mathbb{H}$ a j-invariant of a complex lattice, i.e. a complex lattice up to homothety. We will refer to this function as the *j-function.* Let us define the set

$$\mathcal{F} = \{\omega \in \mathbb{H} \colon \Re(\omega) \in [-1/2, 1/2], |\omega| \geq 1 \text{ and, if } \Re(\omega) > 0, \text{ then } |\omega| > 1\}.$$

One can show the following theorem.

Theorem 2.1.13. *The restriction of the j-function to $\mathcal{F}$ defines a bijection from $\mathcal{F}$ to $\mathbb{C}$.*

Proof. See [1, Theorem 2.7]. □

This theorem allows us to complete our picture with the following corollary, known as *Uniformization Theorem.*

Corollary 2.1.14 (Uniformization Theorem)**.** *For every elliptic curve $E/\mathbb{C}$, there exists a complex lattice Λ such that $E = E_\Lambda$.*

Proof. By Theorem 2.1.13, there exists $\omega \in \mathcal{F}$ such that $j(\omega) = j(E)$. Setting $\Lambda = [1, \omega]$, we can conclude. □

If we want to summarize, elliptic curves over the complex field up to isomorphism are in one-to-one correspondence with complex lattices up to homothety.

We conclude this section with a lemma that will be useful in the following.

Lemma 2.1.15. *Let $\Lambda = [1, \tau]$ be a complex lattice with $\tau \in \mathbb{H}$ and let p be a prime number. The only cyclic sublattices of Λ of index p are the lattice $[1, p\tau]$ and the lattices $[p, \tau + k]$, for $0 \leq k < p$.*

Proof. The lattices $[1, p\tau]$ and $[p, \tau + k]$ are clearly index p sublattices of Λ and they must be cyclic, since p is prime. Conversely, any sublattice $\Lambda' \subseteq \Lambda$ can be written as $[d, a\tau + k]$, where d is the smallest positive integer in Λ'. The index of Λ' in Λ is ad. Since p is prime and we are looking for index p sublattices, we must either have $d = 1$ and $a = p$ or $d = p$ and $a = 1$. The first case corresponds to the lattice $[1, p\tau]$, while the second case corresponds to the lattices $[p, \tau + k]$, and we may assume that $0 \leq k < p$. □

2.2 Isogenies and Lattices Inclusions

Recall that an holomorphic map between complex manifolds is usually defined locally on charts. However, the situation is much simpler in the case of complex tori, since they only have one global chart. Let Λ_1, Λ_2 be two complex lattices and $\mathbb{C}/\Lambda_1, \mathbb{C}/\Lambda_2$ their corresponding complex tori. Their complex structure can be described by one global chart, i.e. the quotient map

$$p_i \colon \mathbb{C} \longrightarrow \mathbb{C}/\Lambda_i,$$

for $i = 1, 2$. A map $\varphi \colon \mathbb{C}/\Lambda_1 \longrightarrow \mathbb{C}/\Lambda_2$ between complex tori is holomorphic if it is induced by an holomorphic map $f \colon \mathbb{C} \longrightarrow \mathbb{C}$ that makes the diagram commute:

$$\begin{tikzcd}
\mathbb{C} \arrow[r, "f"] \arrow[d, "p_1"] & \mathbb{C} \arrow[d, "p_2"] \\
\mathbb{C}/\Lambda_1 \arrow[r, "\varphi"] & \mathbb{C}/\Lambda_2.
\end{tikzcd}$$

We will call *morphisms of complex tori* the holomorphic maps between complex tori. Every $\alpha \in \mathbb{C}$ determines a holomorphic map $z \longmapsto \alpha z$, which is an endomorphism of $\mathbb{C}$ as an additive group. If $\alpha\Lambda_1 \subseteq \Lambda_2$, this map induces a group morphism

$$\varphi_\alpha \colon \mathbb{C}/\Lambda_1 \longrightarrow \mathbb{C}/\Lambda_2, \quad z + \Lambda_1 \longmapsto \alpha z + \Lambda_2,$$

and one can verify that this is a morphism of complex tori.

Lemma 2.2.1. *Let $\varphi \colon \mathbb{C}/\Lambda_1 \longrightarrow \mathbb{C}/\Lambda_2$ be a morphism of complex tori with $\varphi(0) = 0$. There is a unique $\alpha \in \mathbb{C}$ for which $\varphi = \varphi_\alpha$.*

Proof. See [33, Theorem VI.4.1]. □

The next theorem shows the precise relation between lattice inclusions and morphisms of complex tori.

Theorem 2.2.2. *Let Λ_1, Λ_2 be two complex lattices. The map*

$$\{\alpha \in \mathbb{C} \colon \alpha\Lambda_1 \subseteq \Lambda_2\} \to \{\textit{morphisms } \varphi \colon \mathbb{C}/\Lambda_1 \longrightarrow \mathbb{C}/\Lambda_2\}, \quad \alpha \mapsto \varphi_\alpha$$

is an isomorphism of additive groups. If $\Lambda_1 = \Lambda_2$, it is an isomorphism of commutative rings.

Proof. By Lemma 2.2.1, we know that the above map is a bijection, so we just need to prove it is a morphism. Let $\alpha, \beta \in \mathbb{C}$ be such that $\alpha\Lambda_1 \subseteq \Lambda_2$ and $\beta\Lambda_1 \subseteq \Lambda_2$. By commutativity of the diagram, we have

$$\varphi_{\alpha+\beta}(p_1(z)) = p_2((\alpha\beta z)) = p_2(\alpha z) + p_2(\beta z) = \varphi_\alpha(p_1(z)) + \varphi_\beta(p_1(z)) = (\varphi_\alpha + \varphi_\beta)(p_1(z)),$$

as we wanted to prove.

If $\Lambda_1 = \Lambda_2$, letting $p = p_1 = p_2$, we also have

$$\varphi_{\alpha\beta}(p(z)) = p(\alpha\beta z) = \varphi_\alpha(\beta z) = (\varphi_\alpha \circ \varphi_\beta)(p(z)),$$

which shows that the above map is also a ring morphism. □

In the following, we will often identify morphisms of complex tori with the set

$$\{\alpha \in \mathbb{C} \colon \alpha\Lambda_1 \subseteq \Lambda_2\}.$$

As one may expect, this set in one-to-one correspondence also with with isogenies between the elliptic curves E_{Λ_1} and E_{Λ_2}.

Theorem 2.2.3. *Let Λ_1, Λ_2 be two complex lattices and $E_1 = E_{\Lambda_1}$, $E_2 = E_{\Lambda_2}$ the corresponding elliptic curves. Let $\Phi_i \colon \mathbb{C}/\Lambda_i \longrightarrow E_i$ be the isomorphisms as in Theorem 2.1.7 and $\alpha \in \mathbb{C}$. We have $\alpha\Lambda_1 \subseteq \Lambda_2$ if and only if there exists a uniquely determined $\phi_\alpha \in \mathrm{Hom}(E_1, E_2)$ such that the following diagram commutes:*

$$\begin{tikzcd}
\mathbb{C}/\Lambda_1 \arrow[r, "\Phi_1"] \arrow[d, "\alpha"] & E_1 \arrow[d, "\phi_\alpha"] \\
\mathbb{C}/\Lambda_2 \arrow[r, "\Phi_2"] & E_2.
\end{tikzcd}$$

Conversely, for every morphism $\phi \in \operatorname{Hom}(E_1, E_2)$ there exists a unique $\alpha_\phi \in \mathbb{C}$ such that $\alpha_\phi \Lambda_1 \subseteq \Lambda_2$. The maps $\phi \longmapsto \alpha_\phi$ and $\alpha \longmapsto \phi_\alpha$ are inverse isomorphisms between $\operatorname{Hom}(E_1, E_2)$ *and* $\{\alpha \in \mathbb{C} \colon \alpha\Lambda_1 \subseteq \Lambda_2\}$.

Proof. See [33, Theorem VI.4.1]. Here, we prove that the isomorphism $\Psi : \phi \longmapsto \alpha_\phi$ is a group morphism and that it is the inverse of $\alpha \longmapsto \phi_\alpha$. We have $\Psi(0) = 0$ by commutativity of the diagram, and for all $\phi_1, \phi_2 \in \operatorname{Hom}(E_1, E_1)$

$$\Psi(\phi_1 + \phi_2) = \Phi_2^{-1} \circ (\phi_1 + \phi_2) \circ \Phi_1 = \Phi_2^{-1} \circ \phi_1 \circ \Phi_1 + \Phi_2^{-1} \circ \phi_2 \circ \Phi_1 = \Psi(\phi_1) + \Psi(\phi_2),$$

again exploiting the commutativity of the diagram. This implies that Ψ is a group isomorphism. Let now $\psi \in \operatorname{Hom}(E_1, E_2)$. The map $\Phi_2^{-1} \circ \psi \circ \Phi_1$ is a morphism of complex tori and, by Theorem 2.2.2, it is induced by a uniquely determined α_ψ satisfying $\alpha_\psi \Lambda_1 \subseteq \Lambda_2$. There exists a uniquely determinated ϕ_{α_ψ} that makes the diagram commute. Hence, we get $\phi_{\alpha_\psi} = \psi$, which proves that the two isomorphisms are inverses to each other.

□

Corollary 2.2.4. *Let Λ be a complex lattice and $E = E_\Lambda$ be its corresponding elliptic curve. The maps defined in the previous theorem $\alpha \longmapsto \phi_\alpha$ and $\phi \longmapsto \alpha_\phi$ are inverse ring isomorphisms between $\{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\}$ and* $\operatorname{End}(E)$. *The involution map $\phi \longmapsto \widehat{\phi}$ in* $\operatorname{End}(E)$ *corresponds to complex conjugation in $\{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\}$ and we have* $\operatorname{Tr}(\phi_\alpha) = \alpha + \overline{\alpha}$ *and* $\deg \phi_\alpha = \alpha\overline{\alpha}$.

Proof. Let $\Psi \colon \mathbb{C} \longrightarrow E$ be the isomorphism of Theorem 2.1.7 and

$$\Psi \colon \operatorname{End}(E) \longrightarrow \{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\}$$

be the isomorphism of groups that maps every $\phi \in \operatorname{End}(E)$ into α_ϕ. We want to prove that Ψ is also a ring morphism. Let $\phi_1, \phi_2 \in \operatorname{End}(E)$. We have

$$\Psi(\phi_1 \circ \phi_2) = \Phi_2^{-1} \circ (\phi_1 \circ \phi_2) \circ \Phi_1 = (\Phi_2^{-1} \circ \phi_1 \circ \Phi_1) \circ (\Phi_2^{-1} \circ \phi_2 \circ \Phi_1) = \Psi(\phi_1) \circ \Psi(\phi_2),$$

by the commutativity of the diagram. This implies that Ψ is a ring isomorphism.

Let $\phi \in \operatorname{End}(E)$. Identifying $\operatorname{End}(E)$ with the set $\{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\}$, the complex number $\alpha = \alpha_\phi$ satisfies the equation

$$x^2 - \operatorname{Tr}(\phi)x + \operatorname{N}(\phi) = 0,$$

which has integer coefficients and discriminant $\operatorname{Tr}(\phi)^2 - 4\operatorname{N}(\phi) \leq 0$. Hence, either $\alpha \in \mathbb{Z}$ or α is an algebraic integer in a quadratic imaginary number field. In both cases, its conjugate $\overline{\alpha}$ satisfies the same equation. Thus, we have $\alpha + \overline{\alpha} = \operatorname{Tr}(\phi)$ and $\alpha\overline{\alpha} = \operatorname{N}(\phi) = \deg \phi = \widehat{\phi}\phi$, which implies that complex conjugation corresponds to the dual map involution, i.e. $\Psi(\widehat{\phi}) = \overline{\alpha}$. This is because $\operatorname{End}(E)$ has no zero divisor, by Corollary 1.2.15, so we can apply the cancellation law. □

Corollary 2.2.5. *Let $E/\mathbb{C}$ be an elliptic curve. Then its endomorphism ring is commutative, and thus it is isomorphic either to $\mathbb{Z}$ or an order in an imaginary quadratic number field.*

Remark 2.2.6. The above discussion justifies the origin of the term complex multiplication. When the endomorphism ring of E_Λ strictly contains $\mathbb{Z}$, all the extra morphisms correspond to the morphism of complex tori given by the multiplication by α, for some α algebraic integer lying in an imaginary quadratic number field. On the other hand, if the endomorphism ring of E_Λ is exactly $\mathbb{Z}$, then it consists only of multiplication by m maps.

2.3 Class Group Action

In this section, we are going to define the class group of an order $\mathcal{O}$ in an imaginary quadratic number field and its action on the set of elliptic curves with complex multiplication by $\mathcal{O}$. However, we first need to recall some ideas from number theory.

Recall that in the previous section we proved that we have three isomorphic representations of the endomorphism ring of a complex elliptic curve, i.e.

$$\operatorname{End}(E_\Lambda) \simeq \{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\} \simeq \{\text{endomorphisms of } \mathbb{C}/\Lambda\}. \tag{2.2}$$

In the following, we will identify them.

Let $\mathcal{O}$ be an order in the quadratic imaginary number field L. Notice that $\mathcal{O}$ is a complex lattice with respect to its additive structure. This can be proved using Theorem 1.5.12 and the known structure of $\mathcal{O}_L$, the ring of integers of L. In the previous section, we proved that every elliptic curve E over the complex numbers arise from a complex lattice Λ. Suppose E has complex multiplication by $\mathcal{O}$. We would like to understand how the lattices $\mathcal{O}$ and Λ are related. In particular, we want to find every complex lattice Λ such that $\operatorname{End}(E_\Lambda) = \mathcal{O}$.

Suppose $\Lambda = \mathcal{O}$. Let $\alpha \in \operatorname{End}(E_\mathcal{O})$. Exploiting identification (2.2), we have $\alpha\mathcal{O} \subseteq \mathcal{O}$ and so $\alpha \in \mathcal{O}$, since $\mathcal{O}$ is a ring with identity. Conversely, if $\alpha \in \mathcal{O}$, we have $\alpha\mathcal{O} \subseteq \mathcal{O}$ trivially, since $\mathcal{O}$ is a ring. Therefore, we have $\alpha \in \operatorname{End}(E_\mathcal{O})$. This proves that $\operatorname{End}(E_\mathcal{O}) = \mathcal{O}$. Notice that we should expect that the same holds for lattices homothetic to $\mathcal{O}$, since the endomorphism ring is invariant under isomorphism and homothetic lattices produce isomorphic elliptic curves. This is indeed the case, since the set $\{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\}$ does not change if we substitute Λ with $z\Lambda$ for some $z \in \mathbb{C}$.

Since we are working up to homothety and isomorphism, we can assume that $\Lambda = [1, \tau]$. Moreover, from basic number theory, we know that $\mathcal{O}_L = \mathbb{Z}[\mu]$, and there are two possibilities for μ, namely

$$\mu = \begin{cases} \sqrt{d_L}, \\ \dfrac{d_L + \sqrt{d_L}}{2}, \end{cases}$$

depending on the specific extension. Furthermore, by Theorem 1.5.12, we have that

$$\mathcal{O} = \mathbb{Z} + f\mathcal{O}_L,$$

for some $f \in \mathbb{N}$, called *conductor* of the order $\mathcal{O}$. Thus, letting $\omega = f\mu$, we can write $\mathcal{O} = [1, \omega]$.

Suppose $\mathrm{End}(E_\Lambda) = \mathcal{O}$. Then, we must have $\omega\Lambda \subseteq \Lambda$, which implies

$$\omega \cdot 1 = \omega \in \Lambda.$$

Hence, there exist $m, n \in \mathbb{Z}$ such that $\omega = m + n\tau$. This implies that $n\Lambda = [n, \omega - m] = [n, \omega]$, i.e. Λ is homothetic to a sublattice of index n of $\mathcal{O}$. The sublattice $n\Lambda$ must be closed under multiplication by $\mathcal{O}$, so Λ is an $\mathcal{O}$-ideal. However, two sublattices $[m, \omega]$ and $[m, \omega]$ can be homothetic even if $n \neq m$. Moreover, not every $\mathcal{O}$-ideal produces an elliptic curve whose endomorphism ring is exactly $\mathcal{O}$. Let $\Lambda \subseteq \mathcal{O}$ be an ideal. The set

$$\mathcal{O}(\Lambda) = \{\alpha \in \mathbb{C} \colon \alpha\Lambda \subseteq \Lambda\} = \{\alpha \in \ L \colon \alpha\Lambda \subseteq \Lambda\}$$

is an order in L. We have $\mathcal{O} \subseteq \mathcal{O}(\Lambda) = \mathrm{End}(E_\Lambda)$, since Λ is closed under multiplication by $\mathcal{O}$, which implies that $\mathrm{End}^A(E_\Lambda) = L$. However, it is not necessarily true that $\mathcal{O}(\Lambda) = \mathcal{O}$.

Remark 2.3.1. Heuristically, there are many more complex elliptic curves without complex multiplication, which means that their endomorphism ring is simply $\mathbb{Z}$, than complex elliptic curves with complex multiplication. Indeed, as we have just observed, curves of the latter type are of the form E_Λ, with Λ (homothetic to) either a quadratic imaginary order or a proper ideal of a quadratic imaginary order. Hence, all other homotety classes of lattices correspond to curves whose endomorphism ring is isomorphic to $\mathbb{Z}$. For example, if we take $\Lambda = [1, \sqrt[3]{2}]$, then E_Λ has not complex multiplication.

Definition 2.3.2. Let $\mathcal{O}$ be an order in a quadratic imaginary number field L and let $\Lambda \subseteq \mathcal{O}$ be an ideal. We say that Λ is a *proper $\mathcal{O}$-ideal* if $\mathcal{O}(\Lambda) = \mathcal{O}$.

Since we want to study lattices up to homothety, we should consider $\mathcal{O}$-ideals up to homothety. Assume we have $\mathfrak{b} = z\mathfrak{a}$ for some $z \in \mathbb{C}^*$, with $\mathfrak{a}, \mathfrak{b}$ $\mathcal{O}$-ideals. We can always suppose that $z = \alpha/\beta$, for some $\alpha, \beta \in \mathcal{O}$. Indeed, if $\mathfrak{a} = [\omega_1, \omega_2]$ we can take $\alpha = z\omega_1$ and $\beta = \omega_1$. Hence, homothetic ideals always satisfy a relation $\alpha\mathfrak{a} = \beta\mathfrak{b}$ for some $\alpha, \beta \in \mathcal{O}$.

Definition 2.3.3. Let $\mathcal{O}$ be an order in an imaginary quadratic number field and $\mathfrak{a}, \mathfrak{b} \subseteq \mathcal{O}$ two ideals. We say that $\mathfrak{a}$ and $\mathfrak{b}$ are *equivalent ideals* if they are homothetic as lattices, i.e. there exist $\alpha, \beta \in \mathcal{O}$ such that $\alpha\mathfrak{a} = \beta\mathfrak{b}$.

The relation of being equivalent for $\mathcal{O}$-ideals is an equivalence relation.

Definition 2.3.4. Let $\mathcal{O}$ be an order in a quadratic imaginary number field. The *ideal class group* $\mathrm{cl}(\mathcal{O})$ is the multiplicative group of equivalence classes of proper $\mathcal{O}$-ideal.

It is not clear *a priori* that the ideal class group is actually a group, and we are going to prove it.

From our previous discussion, we can draw the following conclusion.

Theorem 2.3.5. *Let $\mathcal{O}$ be an order in a quadratic imaginary number field. There is a one-to-one correspondence between elements of the ideal class group* $\mathrm{cl}(\mathcal{O})$ *and homothety classes of lattices* $\Lambda \subseteq \mathbb{C}$ *for which* $\mathrm{End}(E_\Lambda) \simeq \mathcal{O}$.

It is well-known that the class group is finite. Its cardinality is called *class number* and we will denote it by $\mathrm{h}(\mathcal{O})$.

Now, we want to gain a better understanding of what proper $\mathcal{O}$-ideals are. We recall the definition of fractional and invertible ideals.

Definition 2.3.6. Let $\mathcal{O}$ be an integral domain with fraction field L. For any $\alpha \in L$ and $\mathfrak{a} \subseteq \mathcal{O}$ ideal, the $\mathcal{O}$-module $\mathfrak{b} = \alpha\mathfrak{a}$ is called *fractional ideal.*

If it is not clear from the context, we will call $\mathcal{O}$-ideals integral ideals. Fractional ideals lying in $\mathcal{O}$ are integral ideals, and every $\mathcal{O}$-ideal is a fractional ideal. Fractional ideals multiply in the obvious way and $\mathcal{O}$ is a fractional ideal too, acting as the multiplicative identity. If $\mathcal{O}$ is an order in a number field, we can always write nonzero fractional ideals as $\mathfrak{b} = \frac{1}{b}\mathfrak{a}$ for some $b \in \mathbb{Z}$, $b > 0$.

Definition 2.3.7. Fractional ideals $\mathfrak{b}$ for which there exists another fractional ideal $\mathfrak{b}^{-1}$ such that $\mathfrak{b}\mathfrak{b}^{-1} = \mathcal{O}$ are said *invertible ideals.*

Not every fractional ideal is invertible. For example, the zero ideal never is. In our case, if $\mathcal{O}$ is strictly contained in the ring of integers, then there are nonzero non-invertible fractional $\mathcal{O}$-ideals. The set of invertible ideals forms a group, called *ideal group.*

Let L be a number field. Recall that the norm $\mathrm{N}_{L/\mathbb{Q}}$ and the trace $\mathrm{Tr}_{L/\mathbb{Q}}$ of any element $\alpha \in L$ are defined as the determinant and the trace of the multiplication by α map, which is a linear invertible transformation. If L is quadratic and imaginary, we have

$$\mathrm{N}_{L/\mathbb{Q}}(\alpha) = \alpha\overline{\alpha}, \quad \mathrm{Tr}_{L/\mathbb{Q}}(\alpha) = \alpha + \overline{\alpha},$$

so it coincides with the definition of norm and trace given in the previous chapter. If the extension is understood, we will omit it in the notation.

Definition 2.3.8. Let $\mathcal{O}$ be an order in any number field L and let $\mathfrak{a}$ be a nonzero $\mathcal{O}$-ideal. The (*absolute*) *norm* of the ideal $\mathfrak{a}$ is defined as

$$\mathfrak{N}(\mathfrak{a}) = [\mathcal{O} \colon \mathfrak{a}] = \#\frac{\mathcal{O}}{\mathfrak{a}} \in \mathbb{Z}_{>0}.$$

One can also interpret $\mathfrak{N}(\mathfrak{a})$ as the ratio of the volumes of the fundamental domains of $\mathcal{O}$ and $\mathfrak{a}$, viewed as complex lattices.

In the following lemma, we will give two useful properties of the absolute norm of ideals.

Lemma 2.3.9. *Let $\mathcal{O}$ be an order in the number field L, $\alpha \in \mathcal{O}$ a nonzero element, and $\mathfrak{a}$ a nonzero $\mathcal{O}$-ideal.*

(1) *We have $\mathfrak{N}((\alpha)) = |\,\mathrm{N}(\alpha)|$, where (α) denotes the principal ideal generated by α.*

(2) *We have $\mathfrak{N}(\alpha\mathfrak{a}) = |\,\mathrm{N}(\alpha)|\mathfrak{N}(\mathfrak{a})$.*

Proof. See [10, Lemma 7.14]. □

Definition 2.3.10. Let $\mathfrak{b} = \frac{1}{b}\mathfrak{a}$ be a non zero fractional ideal in an order $\mathcal{O}$ in a number field L, with $b \in \mathbb{Z}, b > 0$. The (absolute) norm of $\mathfrak{b}$ is defined as

$$\mathfrak{N}(\mathfrak{b}) = \frac{\mathfrak{N}(\mathfrak{a})}{\mathrm{N}(b)}.$$

Lemma 2.3.9 ensures that the above definition is well defined and does not depend on the choice of $\mathfrak{a}$ and b. Moreover, when $\mathfrak{b}$ is an integral ideal, we have $b = 1$ and the definitions agree.

Let $\mathcal{O}$ be an order in a quadratic imaginary number field L. We extend our definition of proper ideals to fractional ones. Hence, we denote

$$\mathcal{O}(\mathfrak{b}) = \{\alpha \colon \alpha\mathfrak{b} \subseteq \mathfrak{b}\}$$

and say that $\mathfrak{b}$ is proper if $\mathcal{O}(\mathfrak{b}) = \mathfrak{b}$.

Lemma 2.3.11. *Let $\mathcal{O}$ be an order in an imaginary quadratic number field L, $\mathfrak{a}$ be a nonzero $\mathcal{O}$-ideal and $\mathfrak{b} = \alpha\mathfrak{a}$, for $\alpha \in L$. Then, $\mathfrak{b} = [a, b]$ as a lattice, for some $a, b \in L$, which means it is a rank 2 $\mathbb{Z}$-module. Moreover, $\mathfrak{a}$ is proper if and only if $\mathfrak{b}$ is proper, and $\mathfrak{a}$ is invertible if and only if $\mathfrak{b}$ is invertible.*

Proof. Any fractional $\mathcal{O}$-ideal $\mathfrak{b}$ is a $\mathbb{Z}$-submodule of $\mathcal{O}$. Since $\mathcal{O}$ is free, $\mathfrak{b}$ is free, too. Since the rank of $\mathcal{O}$ is two, also the rank of $\mathfrak{b}$ is two because we have

$$\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q} \simeq L \quad \text{and} \quad \mathfrak{b} \otimes_{\mathbb{Z}} \mathbb{Q} \simeq L.$$

We have

$$\{\beta \colon \beta\mathfrak{b} \subseteq \mathfrak{b}\} = \{\beta \colon \beta\mathfrak{a} \subseteq \mathfrak{a}\}.$$

This proves the second statement. If $\mathfrak{a}$ is invertible, we can define $\mathfrak{b}^{-1} = \alpha^{-1}\mathfrak{a}^{-1}$ and verify that this is the inverse of $\mathfrak{b}$. If $\mathfrak{b}$ is invertible, then $\mathfrak{a}^{-1} = \alpha\mathfrak{b}^{-1}$ is an inverse for $\mathfrak{a}$. □

The following theorem unifies the notions of invertible and proper fractional ideals.

Theorem 2.3.12. *Let $\mathcal{O}$ be an order in an imaginary quadratic number field, and $\mathfrak{a} = [a, b]$ be a fractional $\mathcal{O}$-ideal. Then, $\mathfrak{a}$ is proper if and only if it is invertible. If $\mathfrak{a}$ is invertible, we have $(\mathfrak{N}(\mathfrak{a})) = \overline{\mathfrak{a}}\mathfrak{a}$, where $\overline{\mathfrak{a}} = [\overline{a}, \overline{b}]$ and $(\mathfrak{N}(\mathfrak{a}))$ is the principal ideal generated by the absolute norm of $\mathfrak{a}$. The inverse of $\mathfrak{a}$ is the fractional ideal $\mathfrak{a}^{-1} = \frac{1}{\mathfrak{N}(\mathfrak{a})}\overline{\mathfrak{a}}$.*

Proof. See [10, Proposition 7.4]. □

Corollary 2.3.13. *The ideal class group $\mathrm{cl}(\mathcal{O})$ of an order $\mathcal{O}$ in an imaginary quadratic number field L is the group of invertible fractional ideals of $\mathcal{O}$ modulo its subgroup of principal fractional ideals. In particular, $\mathrm{cl}(\mathcal{O})$ is a group.*

Proof. Let G be the group of invertible fractional ideals and H its subgroup of principal fractional ideals. Each invertible fractional ideal $\mathfrak{b} = \frac{1}{b}\mathfrak{a}$ is the product of an invertible principal fractional ideal $(\frac{1}{b})$ and an invertible ideal $\mathfrak{a}$, by Lemma 2.3.11. This implies that G/H consists of equivalence classes $\mathfrak{a}H$, where $\mathfrak{a}$ is an invertible (equivalently, proper) ideal. Each nonzero principal fractional ideal is invertible, since $(\alpha)^{-1} = (\alpha^{-1})$. Hence, H contains every nonzero principal fractional ideal and for any two invertible/proper ideals $\mathfrak{a}, \mathfrak{b}$, $\mathfrak{a}$ and $\mathfrak{b}$ are in the same equivalence class if and only if $\mathfrak{a}H = \mathfrak{b}H$. This implies that $\mathrm{cl}(\mathcal{O}) = G/H$. □

The ring of integers of a number field is the maximal order and it has many good properties. For example, all ideals are invertible and all ideals can be written as the product of prime ideals in a unique way. In general, if we consider non-maximal orders, we cannot hope for a unique factorization of ideals and not all ideals are invertible, as we have seen in the proof of Theorem 2.3.15. However, if we consider ideals prime to the conductor, these two properties still hold true in general orders. For example, we have the following lemma.

Lemma 2.3.14. *Let $\mathcal{O}$ be an imaginary quadratic order. A nonzero prime $\mathcal{O}$-ideal is invertible if and only if it is prime to the conductor of $\mathcal{O}$.*

Proof. See [8, Theorem 6.1]. □

The next two propositions will be useful in Chapter 3.

Proposition 2.3.15. *Let $\mathcal{O}$ be an order of discriminant d and conductor f in the imaginary quadratic number field L and let ℓ be a prime number. If ℓ divides f of $\mathcal{O}$ there are no proper $\mathcal{O}$-ideals of norm ℓ. Otherwise, there are exactly $1 + (\frac{d}{\ell})$.*

Proof. If ℓ does not divide the conductor of $\mathcal{O}$, then the theorem follows combining the classical number theoretic results [24, Theorem 27] and [10, Proposition 7.22].

If ℓ divides the conductor, then there is no invertible prime ideal of norm ℓ, by Lemma 2.3.14. By Theorem 2.3.12, invertible ideals are proper ideals, and so there is no proper $\mathcal{O}$ ideal of norm ℓ. □

Proposition 2.3.16. *Let $\mathcal{O}$ be an imaginary quadratic order. For each ideal class $[\mathfrak{a}]$ in $\mathrm{cl}(\mathcal{O})$ there are infinitely many prime numbers that are norms of ideals in the class $[\mathfrak{a}]$.*

Proof. This follows from [10, Theorem 7.7] and [10, Theorem 9.12]. □

Corollary 2.3.17. *Let $\mathcal{O}$ be an order in an imaginary quadratic number field and $\mathfrak{a}, \mathfrak{b}$ be invertible/proper fractional ideals. Then $\mathfrak{N}(\mathfrak{a}\mathfrak{b}) = \mathfrak{a}\mathfrak{b}$, i.e. the absolute norm is multiplicative for proper/invertible fractional ideals.*

Proof. By Proposition 2.3.9, it is enough to consider the case where both $\mathfrak{a}$ and $\mathfrak{b}$ are invertible integral ideals. We have

$$(\mathfrak{N}(\mathfrak{a}\mathfrak{b})) = \mathfrak{a}\mathfrak{b}\overline{\mathfrak{a}\mathfrak{b}} = \mathfrak{a}\overline{\mathfrak{a}}\mathfrak{b}\overline{\mathfrak{b}} = (\mathfrak{N}(\mathfrak{a}))(\mathfrak{N}(\mathfrak{b})).$$

It follows that $\mathfrak{N}(\mathfrak{a}\mathfrak{b}) = \mathfrak{N}(\mathfrak{a})\mathfrak{N}(\mathfrak{b})$, since $\mathfrak{N}(\mathfrak{a}\mathfrak{b}), \mathfrak{N}(\mathfrak{a}), \mathfrak{N}(\mathfrak{b}) \in \mathbb{Z}$ are all strictly positive integers. □

We will denote the set of all j-invariants of elliptic curves over K with endomorphism ring $\mathcal{O}$ by $\mathrm{Ell}_{\mathcal{O}}(K)$. Let $\mathcal{O}$ be an order in an imaginary quadratic number field L. We are now ready to introduce the class group action of $\mathcal{O}$ on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$.

Given any $E \in \mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$, there exists a proper $\mathcal{O}$-ideal $\mathfrak{b}$ such that $E = E_{\mathfrak{b}}$. If $\mathfrak{a}$ is any proper/invertible ideal, we define $\mathfrak{a}E_{\mathfrak{b}} = E_{\mathfrak{a}^{-1}\mathfrak{b}}$. Notice that for any $\alpha \in L$, we have $(\alpha\mathfrak{a})E_{\mathfrak{b}} = \mathfrak{a}E_{\mathfrak{b}}$, since the lattices $(\alpha\mathfrak{a})^{-1}\mathfrak{b} = \alpha\mathfrak{a}^{-1}\mathfrak{b}$ and $\mathfrak{a}^{-1}\mathfrak{b}$ are homothetic. This induces an action of $\mathrm{cl}(\mathcal{O})$ on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ defined as

$$\mathrm{cl}(\mathcal{O}) \times \mathrm{Ell}_{\mathcal{O}}(\mathbb{C}) \longrightarrow \mathrm{Ell}_{\mathcal{O}}(\mathbb{C}), \quad (\mathfrak{a}, j(E_{\mathfrak{b}})) = j(E_{\mathfrak{a}^{-1}\mathfrak{b}}).$$

We will write equivalently

$$[\mathfrak{a}]E_{\mathfrak{b}} \quad \text{or} \quad \mathfrak{a}E_{\mathfrak{b}} \quad \text{or} \quad [\mathfrak{a}]j(\mathfrak{b}) \quad \text{or} \quad \mathfrak{a}j(\mathfrak{b})$$

to denote the action of the class of the ideal $\mathfrak{a}$ on the curve of j-invariant $j(E_{\mathfrak{b}})$. Clearly, the identity of the class group acts trivially on $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ and we have

$$\mathfrak{a}(\mathfrak{b}E_{\mathfrak{c}}) = \mathfrak{a}E_{\mathfrak{b}^{-1}\mathfrak{c}} = E_{(\mathfrak{b}\mathfrak{a})^{-1}\mathfrak{c}} = (\mathfrak{a}\mathfrak{b})E_{\mathfrak{c}}.$$

Thus this definition actually provides a group action.

For any proper ideals $\mathfrak{a}$ and $\mathfrak{b}$, we have $[\mathfrak{a}]j(\mathfrak{b}) = j(\mathfrak{b})$ if and only if $\mathfrak{a}^{-1}\mathfrak{b}$ and $\mathfrak{b}$ are homothetic lattices, i.e. $\mathfrak{a}\mathfrak{b} = \alpha\mathfrak{b}$ for some $\alpha \in L$. This implies that $\mathfrak{a} = \alpha\mathcal{O}$ is principal. Thus, the action is *free*, i.e. the identity of $\mathrm{cl}(\mathcal{O})$ is the only element that fixes all elements in $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$.

By Theorem 2.3.5, we have that $\mathrm{cl}(\mathcal{O})$ and $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ have the same cardinality. Therefore, the action must also be *transitive*, i.e. for any fixed $j \in \mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$, its $\mathrm{cl}(\mathcal{O})$-orbit covers the whole $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$.

The $\mathrm{cl}(\mathcal{O})$-action is the fundamental tool to build many cryptographic primitives, and we will see an application in Chapter 4.

Let $\phi\colon E_{\Lambda_1} \longrightarrow E_{\Lambda_2}$ be an isogeny of elliptic curves over $\mathbb{C}$. By Theorem 2.2.3, there exists a unique α with $\alpha\Lambda_1 \subseteq \Lambda_2$ such that the diagram

$$\begin{tikzcd}
\mathbb{C}/\Lambda_1 \arrow[r, "\Phi_1"] \arrow[d, "\alpha"] & E_{\Lambda_1} \arrow[d, "\phi_\alpha"] \\
\mathbb{C}/\Lambda_2 \arrow[r, "\Phi_2"] & E_{\Lambda_2}
\end{tikzcd}$$

commutes. Since we are interested in curves up to isomorphism and lattices up to homothety, we can replace Λ_1 with $\alpha\Lambda_1$, so that $\alpha = 1$ and the isogeny ϕ is induced by an inclusion $\Lambda_1 \subseteq \Lambda_2$. This corresponds to compose the isogeny with an isomorphism. This proves that, up to isomorphism, every isogeny comes from a lattice inclusion of the shape $\Lambda_1 \subseteq \Lambda_2$. Moreover, we can explicitly determine the kernel of the isogeny ϕ. By commutativity of the diagram, since $\alpha = 1$, the kernel of ϕ is the set

$$\{\Phi_1(z) : z \in \mathbb{C}, \Phi_2(z) = 0\}.$$

The elements of this set are precisely $z \in \Lambda_2$. We have that $\Phi_1(z) = 0$ if and only if $z \in \Lambda_1$, thus

$$\#\ker\phi = [\Lambda_2 \colon \Lambda_1].$$

Since we are working in zero characteristic, ϕ is separable and $\deg\phi = [\Lambda_2 \colon \Lambda_1]$.

Now, suppose that $E_{\Lambda_1}/\mathbb{C}$ has complex multiplication by $\mathcal{O}$. Then Λ_1 is homothetic to a proper ideal $\mathfrak{b}$, so we have $E_{\Lambda_1} = E_{\mathfrak{b}}$. If $\mathfrak{a}$ is an invertible/proper ideal, the lattice inclusion $\mathfrak{b} \subseteq \mathfrak{a}^{-1}\mathfrak{b}$ induces an isogeny

$$\phi_{\mathfrak{a}} \colon E_{\mathfrak{b}} \longrightarrow [\mathfrak{a}]E_{\mathfrak{b}}$$

that corresponds to the action of $\mathfrak{a}$ on $j(E_{\mathfrak{b}})$. If E_{Λ_2} has complex multiplication by $\mathcal{O}$, Λ_2 is homothetic to an invertible/proper ideal $\mathfrak{c}$. We can assume that $\mathfrak{c}$ contains $\mathfrak{b}$. Indeed, if we substitute $\mathfrak{b}$ with $(\mathfrak{N}(\mathfrak{c}))\mathfrak{b}$, which is an homothetic lattice, then $\mathfrak{c} \supseteq \mathfrak{b}$, because we have $(\mathfrak{N}(\mathfrak{c})) = \mathfrak{c}\overline{\mathfrak{c}}$. If we set $\mathfrak{a} = \mathfrak{c}\mathfrak{b}^{-1}$, the isogeny $\phi_{\mathfrak{a}} \colon E_{\mathfrak{b}} \longrightarrow E_{\mathfrak{c}}$ induced by the inclusion $\mathfrak{b} \subseteq \mathfrak{c}$ corresponds to the action of the ideal $\mathfrak{a}$ on $E_{\mathfrak{b}}$. After multiplying $\mathfrak{a}, \mathfrak{b}, \mathfrak{c}$ by integers if necessary, we can assume that $\mathfrak{a}$ is an invertible ideal. Therefore, we proved that all elliptic curves over $\mathbb{C}$ with complex multiplication by $\mathcal{O}$ are isogenous and every isogeny between elliptic curves over $\mathbb{C}$ with complex multiplication by $\mathcal{O}$ is induced by the $\mathrm{cl}(\mathcal{O})$-action.

Definition 2.3.18. Let E/K be an elliptic curve with complex multiplication by an order $\mathcal{O}$ in a quadratic imaginary number field. Let $\mathfrak{a}$ be an $\mathcal{O}$-ideal. The $\mathfrak{a}$-*torsion subgroup* of E is defined as

$$E[\mathfrak{a}] = \{P \in E(\overline{K}) \colon \tau(P) = 0 \text{ for all } \tau \in \mathfrak{a}\}.$$

Notice that in the previous definition we are identifying $\mathcal{O}$ with $\mathrm{End}(E)$ and so $\tau \in \mathfrak{a}$ is viewed as an endomorphism.

Theorem 2.3.19. *Let $\mathcal{O}$ be an order in an imaginary quadratic number field. Let $E/\mathbb{C}$ be an elliptic curve with complex multiplication by $\mathcal{O}$, $\mathfrak{a}$ an invertible $\mathcal{O}$-ideal and $\phi_{\mathfrak{a}}$ be the corresponding isogeny from E to $[\mathfrak{a}]E$. We have* $\ker\phi_{\mathfrak{a}} = E[\mathfrak{a}]$ *and* $\deg\phi_{\mathfrak{a}} = \mathfrak{N}(\mathfrak{a})$.

Proof. By composing with an isomorphism if necessary, we assume without loss for generality that $E = E_{\mathfrak{b}}$, for some invertible $\mathcal{O}$-ideal $\mathfrak{b}$. Let Φ be the isomorphism

$$\mathbb{C}/\mathfrak{b} \longrightarrow E_{\mathfrak{b}}$$

such that $z \longmapsto (\wp(z), \wp'(z))$. Then we have

$$\begin{aligned}\Phi^{-1}(E[\mathfrak{a}]) &= \{z \in \mathbb{C}/\mathfrak{b} \colon \alpha z = 0 \text{ for all } \alpha \in \mathfrak{a}\} = \{z \in \mathbb{C} \colon z\mathfrak{a} \subseteq \mathfrak{b}\}/\mathfrak{b} \\ &= \{z \in \mathbb{C} \colon (z) \subseteq \mathfrak{a}^{-1}\mathfrak{b}\}/\mathfrak{b} = \frac{\mathfrak{a}^{-1}\mathfrak{b}}{\mathfrak{b}} \\ &= \ker\big(\iota \colon \mathbb{C}/\mathfrak{b} \longrightarrow \mathbb{C}/\mathfrak{a}^{-1}\mathfrak{b}\big) = \Phi^{-1}(\ker\phi_{\mathfrak{a}}),\end{aligned}$$

where $\iota \colon \mathbb{C}/\mathfrak{b} \longrightarrow \mathbb{C}/\mathfrak{a}^{-1}\mathfrak{b}$ is the tori morphism induced by inclusion of lattices. This proves the former claim since Φ is a isomorphism. On the other hand, we have

$$\#E[\mathfrak{a}] = [\mathfrak{a}^{-1}\mathfrak{b} \colon \mathfrak{b}] = [\mathfrak{b} \colon \mathfrak{a}\mathfrak{b}] = [\mathcal{O} \colon \mathfrak{a}\mathcal{O}] = [\mathcal{O} \colon \mathfrak{a}] = \mathfrak{N}(\mathfrak{a}),$$

which proves the latter claim. □

Notice that while the curve $[\mathfrak{a}]E$ only depends on the ideal class of $\mathfrak{a}$, the isogeny $\phi_{\mathfrak{a}}$ actually depends on the precise ideal we choose. Taking different ideals in the same class will produce different isogenies from E to $[\mathfrak{a}]E$ with different kernels and degrees.

Corollary 2.3.20. *Let $\mathcal{O}$ be an order in an imaginary quadratic number field. For every $\mathfrak{a}$ invertible $\mathcal{O}$-ideal and every elliptic curve $E/\mathbb{C}$ with complex multiplication by $\mathcal{O}$, the elliptic curves E and $[\mathfrak{a}]E$ are related by an isogeny $\phi_{\mathfrak{a}}$ induced by the $\mathrm{cl}(\mathcal{O})$-action of degree $\mathfrak{N}(\mathfrak{a})$.*

Corollary 2.3.21. *Let $\Lambda_1 \subseteq \Lambda_2$ be an inclusion of complex lattices and $\phi\colon E_1 \longrightarrow E_2$ the corresponding induced isogeny. The sublattice Λ_1 is cyclic if and only if ϕ is a cyclic isogeny.*

Proof. Let

$$\Phi_1\colon \mathbb{C}/\Lambda_1 \longrightarrow E_2$$
$$\Phi_2\colon \mathbb{C}/\Lambda_2 \longrightarrow E_2$$

be the isomorphisms as in Theorem 2.1.7. We have the commutative diagram

$$\begin{tikzcd}
\mathbb{C}/\Lambda_1 \arrow[r, "\Phi_1"] \arrow[d, "\iota"] & E_1 \arrow[d, "\phi"] \\
\mathbb{C}/\Lambda_2 \arrow[r, "\Phi_2"] & E_2,
\end{tikzcd}$$

where ι is the map induced by the lattice inclusion. The kernel of the map ϕ coincides with the group

$$G = \{P \in E_1\colon \Phi_2 \circ \iota \circ \Phi_1^{-1}(P) = O\}.$$

Since both Φ_1 and Φ_2 are isomorphisms, the group G is isomorphic to

$$\{\alpha \in \mathbb{C}/\Lambda_1\colon \alpha = 0 \text{ in } \mathbb{C}/\Lambda_2\} \simeq \Lambda_2/\Lambda_1,$$

from which we can prove the claim. □

2.4 Deuring Theorems

In this section, we are going to build some useful tools that we will exploit in the next chapter. These topics and the proofs involved should deserve at least a whole chapter; therefore, we will not go through their details. We recall the notation $\mathrm{Ell}_{\mathcal{E}}(K)$ for the set of j-invariants of elliptic curves over K with endomorphism ring $\mathcal{E}$.

Definition 2.4.1. Let $\mathcal{O}$ be an order with discriminant d in an imaginary quadratic number field. The *Hilbert class polynomial* is

$$H_d(x) = \prod_{j(E)\in \mathrm{Ell}_{\mathcal{O}}(\mathbb{C})} (x - j(E)).$$

Proposition 2.4.2. *Every Hilbert class polynomial has integer coefficients and has degree* $h(\mathcal{O})$.

Proof. See [10, Theorem 11.1] and [23, Theorem 2.7]. □

This proposition implies that every j-invariant of elliptic curves defined over the complex field is an algebraic integer.

Definition 2.4.3. Let $\mathcal{O}$ be an order with discriminant d in an imaginary quadratic number field. The splitting field of the Hilbert class polynomial $H_d(x)$ over $\mathbb{Q}(\sqrt{d})$ is known as the *ring class field* of the imaginary quadratic order $\mathcal{O}$.

By Proposition 1.1.5, we have that every elliptic curve over the complex field is actually defined over a number field, namely the ring class field.

Now we want to introduce a useful lemma that we will use in the following chapter and that allows us to connect ideals in the ring of integers with ideals prime to the conductor in any other order.

Lemma 2.4.4. *Let* $\mathcal{O}$ *be an order of conductor* f *in the imaginary quadratic number field* L *and let* $\mathcal{O}_L$ *be its ring of integers. If* $\mathfrak{a}$ *is an* $\mathcal{O}_L$*-ideal coprime to* f*, then* $\mathfrak{a} \cap \mathcal{O}$ *is an* $\mathcal{O}$*-ideal prime to* f *of the same norm. If* $\mathfrak{b}$ *is an* $\mathcal{O}$*-ideal prime to* f*, then* $\mathfrak{a}\mathcal{O}_L$ *is an* $\mathcal{O}_L$*-ideal prime to* f *of the same norm.*

Proof. See [10, Proposition 7.20] □

The next lemma will be really useful to establish a relation between two orders involving their discriminants.

Lemma 2.4.5. *Let* L *be a number field. If* $\mathfrak{a} \subseteq \mathfrak{a}'$ *are two nonzero finitely generated* $\mathcal{O}_L$*-submodules, then the index* $[\mathfrak{a}' \colon \mathfrak{a}]$ *is finite and satisfies*

$$d(\mathfrak{a}) = [\mathfrak{a}' \colon \mathfrak{a}]^2 d(\mathfrak{a}'),$$

where $d(\mathfrak{a})$ *and* $d(\mathfrak{a}')$ *are the discriminants of* $\mathfrak{a}$ *and* $\mathfrak{a}'$.

Proof. See [26, Proposition 2.12]. □

Now, we are going to introduce the concept of *reduction* of elliptic curves. We begin with the following field theoretic standard result.

Theorem 2.4.6 (Wedderburn's little Theorem)**.** *Any finite integral domain is a finite field.*

Proof. Let R be a finite integral domain and $a \in R$ be a nonzero element. Consider the multiplication map

$$\begin{aligned} T \colon R &\longrightarrow R \\ x &\longmapsto ax. \end{aligned}$$

Since R has no zero divisor, we have $\ker T = \{0\}$, which means that the map is injective. As R is finite, T is also surjective, and therefore there exists $y \in R$ such that $ay = 1$, i.e. a has a inverse in R. □

Corollary 2.4.7. *Let $\mathcal{O}$ be an order in a number field. Then any prime ideal has norm p^n, for some p prime number and $n \in \mathbb{N}$.*

Proof. Let $\mathfrak{q}$ be a prime ideal in $\mathcal{O}$. We have

$$\mathfrak{N}(\mathfrak{q}) = [\mathcal{O} \colon \mathfrak{q}] = \#\frac{\mathcal{O}}{\mathfrak{q}}.$$

Since the above quotient is a finite integral domain, it is actually a field, and therefore it has p^n elements, where p is a prime number and $n \in \mathbb{N}$. □

Let L be a number field and E be an elliptic curve over L defined by the short Weierstrass equation

$$y^2 = x^3 + Ax + B.$$

If $\mathfrak{q}$ is a prime ideal in $\mathcal{O}_L$, we want to reduce E modulo $\mathfrak{q}$. Suppose that A and B can be written in the form α/β, for some $\alpha, \beta \in \mathcal{O}_L$ with $\beta \notin \mathfrak{q}$. Then we can reduce those coefficients modulo $\mathfrak{q}$, obtaining a curve $\overline{E}$ over $\mathcal{O}_L/\mathfrak{q}$ of equation

$$y^2 = x^3 + [A]x + [B].$$

If, in addition, we get

$$\Delta = [A]^3 - 27[B]^3 \neq 0 \quad \text{in } \mathcal{O}_L/\mathfrak{q},$$

then $\overline{E}$ is an elliptic curve over $\mathcal{O}_L/\mathfrak{q}$. Since $\mathfrak{q}$ is a prime ideal, we have that $\mathcal{O}_L/\mathfrak{q}$ is a finite integral domain, and, by Theorem 2.4.6, it is a finite field. In this case, we call $\overline{E}$ the reduction of E modulo $\mathfrak{q}$, and we say that E has *good reduction* modulo $\mathfrak{q}$. Suppose now that $\operatorname{End}(E) = \mathcal{O}$ is an order in a quadratic imaginary number field. What can we say about $\operatorname{End}(\overline{E})$? If $\phi \in \mathcal{O}$ is a nonzero endomorphism of E, then we can reduce the coefficients of the rational functions defining ϕ modulo the prime ideal $\mathfrak{q}$ to obtain an endomorphism $\overline{\phi} \in \operatorname{End}(\overline{E})$. The endomorphism $\overline{\phi}$ is nonzero, as it must satisfy the equation $x^2 - [\operatorname{Tr}\phi]x + [\deg\phi] = 0$ in $\operatorname{End}(\overline{E})$. Hence, we are left with an injective map

$$\operatorname{End}(E) \longrightarrow \operatorname{End}(\overline{E}).$$

Under some suitable assumptions, this function is actually an isomorphism.

Theorem 2.4.8. *Let E be an elliptic curve over a number field L with complex multiplication by $\mathcal{O}$, a quadratic imaginary order of conductor f. Suppose that the prime number p splits completely in L and $\mathfrak{p}$ is a prime $\mathcal{O}_L$-ideal lying over p. Moreover, assume that p does not divide f and that E has good reduction $\overline{E}$ modulo $\mathfrak{p}$. Then $\overline{E}$ is an ordinary elliptic curve in characteristic p and the map*

$$\operatorname{End}(E) \longrightarrow \operatorname{End}(\overline{E})$$

is an isomorphism.

Proof. See [22, Theorem 13.12]. □

Theorem 2.4.9 (Deuring lifting theorem). *Let $E/\mathbb{F}_q$ be an elliptic curve over a finite field and let $\phi \in \operatorname{End}(E)$ be a nonzero endomorphism. There exists an elliptic curve E^* over a number field L with an endomorphism ϕ^* such that E^* has good reduction modulo a prime ideal $\mathfrak{p}$ of L with residue field $\mathbb{F}_q$ and E and ϕ are their reductions modulo $\mathfrak{p}$.*

Proof. See [22, Theorem 13.14]. □

Chapter 3

Isogeny Graphs

In this chapter, we are going to introduce isogeny graphs and analyse their structure and properties under different constraints.

3.1 Complex Volcanoes

We begin with a lemma which claims that elliptic curve endomorphism algebras are isogeny-invariant.

Lemma 3.1.1. *Let ℓ be a prime number and $\phi\colon E \longrightarrow E'$ be an ℓ-isogeny of elliptic curves defined over K. Then $\operatorname{End}^A(E) \simeq \operatorname{End}^A(E')$.*

Proof. Let $\widehat{\phi}$ be the dual isogeny of ϕ. If $\psi \in \operatorname{End}(E)$, the isogeny $\psi' = \phi \circ \psi \circ \widehat{\phi} \in \operatorname{End}(E')$ satisfies

$$\operatorname{Tr}(\psi') = \phi \circ \psi \circ \widehat{\phi} + \phi \circ \widehat{\psi} \circ \widehat{\phi} = \phi \circ \operatorname{Tr}(\psi) \circ \widehat{\phi} = \operatorname{Tr}(\psi) \circ \widehat{\phi} \circ \phi = \ell \operatorname{Tr}(\psi),$$

$$\operatorname{N}(\psi') = \phi \circ \psi \circ \widehat{\phi} \circ \phi \circ \widehat{\psi} \circ \widehat{\phi} = \phi \circ \psi \circ \ell \circ \widehat{\psi} \circ \widehat{\phi} = \phi \circ \ell \operatorname{N}(\psi) \circ \widehat{\phi} = \ell^2 \operatorname{N}(\psi),$$

because all maps of multiplication by some integer commute with all isogenies. Hence, ψ' is a root of the polynomial $x^2 - \ell \operatorname{Tr}(\psi) x + \ell^2 \operatorname{N}(\psi)$, which implies that $\psi'/\ell \in \operatorname{End}^A(E')$ is a root of $x^2 - \operatorname{Tr}(\psi) x + \operatorname{N}(\psi)$. It follows that the characteristic polynomial of every endomorphism of E has a root in $\operatorname{End}^A(E')$, so that $\operatorname{End}(E) \subseteq \operatorname{End}^A(E')$. Following the same argument in the reverse direction shows that $\operatorname{End}(E') \subseteq \operatorname{End}^A(E)$. □

Definition 3.1.2. Let K be a field and P a set of prime numbers not dividing $\operatorname{char} K$. The *K-rational P-isogeny (multi)graph $G_{K,P}$* is the directed graph whose vertices are isomorphism classes of elliptic curves defined over K and directed edges are ℓ-isogeny between those curves.

If $P = \{\ell\}$, we will denote the K-rational P-isogeny graph simply by $G_{K,\ell}$. If $K = \mathbb{F}_q$, we will denote the $\mathbb{F}_q$-rational P-isogeny graph by $G_{q,P}$.

Usually, since we are considering curves up to isomorphism, each curve is represented by its j-invariant. However, in some special cases, it is preferable to use different kinds of representatives, as we will see in Chapter 4. Moreover, we also want to identify isogenies up to isomorphism, i.e. two isogenies with isomorphic domains are identified if they have the same kernel.

Remark 3.1.3. Notice that the existence of the dual isogeny guarantees that (j_1, j_2) is an edge if and only if (j_2, j_1) is an edge. Moreover, if we exclude the pathological case of the curves with j-invariant 0 and 1728, we can consider $G_{K,P}$ as an undirected graph, because the number of outgoing edges matches the number of ingoing edges. The problem with the j-invariants 0 and 1728 is that they have abundant automorphisms, as we have noticed in Theorem 1.5.13: the corresponding curves $y^2 = x^3 + 1$ and $y^2 = x^3 + x$ have the automorphisms

$$\xi : (x, y) \longmapsto (\xi x, y), \quad i : (x, y) \longmapsto (-x, iy),$$

respectively, where ξ and i have orders 3 and 4 in $\overline{K}$. Here, by an abuse of notation, we denote with the same letter both the endomorphisms and the elements of $\overline{K}$. If $j(E_1) = 0$, $j(E_2) \neq 0$ and ϕ is an ℓ-isogeny between them, the isogenies $\phi, \phi \circ \xi$ and $\phi \circ \xi^2$ will all have different kernels, and so do not identify, while their dual isogenies all have the same kernel. This happens because ξ does not fix all cyclic subgroups of the ℓ-torsion subgroup of E_1, which correspond to all ℓ-isogenies. For example, let us consider $E_1[2] = \{O, (-\xi, 0), (-\xi^2, 0), (-1, 0)\}$. Then, we have

$$\xi(-\xi^j, 0) = (-\xi^{j+1}, 0),$$

so the only fixed cyclic subgroup is the trivial one. The case for $j(E_1) = 1728$ is similar.

Remark 3.1.4. Suppose that $K = \mathbb{F}_q$, for some $q = p^r$. The above definition combined with Lemma 3.1.1 shows that the graph $G_{K,L}$ has at least two connected components, corresponding to supersingular and ordinary curves. If we restrict ourselves to ordinary curves and let P be the set of all prime numbers, by Theorem 1.4.5, $G_{K,P}$ has as many connected components as possible values for the trace of the Frobenius endomorphism.

The next theorem proposes a classification of isogenies of prime degree between curves with complex multiplication by an order in an imaginary quadratic number field.

Theorem 3.1.5. *Let $\phi\colon E \longrightarrow E'$ be an ℓ-isogeny of elliptic curves defined over K. If $L = \mathrm{End}^A(E)$ is an imaginary quadratic number field with discriminant d_L, then $\mathrm{End}(E) = \mathcal{O}$ and $\mathrm{End}(E') = \mathcal{O}'$ are orders in L such that one of the following holds:*

(1) *$\mathcal{O} = \mathcal{O}'$ and we say that ϕ is an* horizontal *isogeny;*

(2) *$[\mathcal{O}\colon \mathcal{O}'] = \ell$ and we say that ϕ is a* descending *isogeny;*

(3) *$[\mathcal{O}'\colon \mathcal{O}] = \ell$ and we say that ϕ is an* ascending *isogeny.*

We refer collectively to the ascending and descending isogenies as vertical *isogenies.*

Proof. Let f, f' be the conductors of $\mathcal{O}$ and $\mathcal{O}'$, respectively. By Lemma 1.5.12, we have

$$\mathcal{O} = \mathbb{Z} + f\mathcal{O}_L, \quad \mathcal{O}' = \mathbb{Z} + f'\mathcal{O}_L.$$

As we have already recalled in the previous chapter, we know that $\mathcal{O}_L = \mathbb{Z}[\mu]$, where μ is

$$\mu = \begin{cases} \sqrt{d_L}, \\ \dfrac{d_L + \sqrt{d_L}}{2}, \end{cases}$$

depending on the specific extension. Letting $\tau = f\mu$ and $\tau' = f'\mu$, we have that $\mathcal{O} = [1, \tau]$ and $\mathcal{O}' = [1, \tau']$ as lattices. Let $\hat{\phi}$ be the dual isogeny of ϕ. We have $\psi' = \phi \circ \tau \circ \hat{\phi} \in \mathcal{O}'$ and $\psi = \hat{\phi} \circ \tau' \circ \phi \in \mathcal{O}$, since they clearly are endomorphisms of E' and E, respectively. Following the computations of Theorem 3.1.1, we can get $\mathrm{Tr}(\psi) = \ell \,\mathrm{Tr}(\tau')$ and $\mathrm{N}(\psi) = \ell^2 \,\mathrm{N}(\tau')$, so that ψ is a root of the polynomial

$$x^2 - \ell \,\mathrm{Tr}(\tau')x + \ell^2 \,\mathrm{N}(\tau').$$

We can explicitly compute the roots of the above polynomial and obtain the endomorphism

$$\psi = \ell \left(\frac{\mathrm{Tr}(\tau') \pm \sqrt{\mathrm{Tr}(\tau')^2 - 4\,\mathrm{N}(\tau')}}{2} \right).$$

The endomorphism between parenthesis is the zero of the polynomial

$$x^2 - \mathrm{Tr}(\tau')x + \mathrm{N}(\tau')$$

and so it coincides with τ' or $\widehat{\tau'}$. Therefore, we can conclude that either $\psi = \ell\tau'$ or $\psi = \ell\widehat{\tau'}$. One could go through the same steps and prove that either $\psi' = \ell\tau$ or $\psi' = \ell\widehat{\tau}$. Since $\psi \in \mathcal{O}$ and $\psi' \in \mathcal{O}'$, in both cases, we have $[1, \ell\tau] \subseteq [1, \tau']$ and $[1, \ell\tau'] \subseteq [1, \tau]$ as lattices. Let us show this just for the inclusion $[1, \ell\tau'] \subseteq [1, \tau]$, since one can work out the details for the other one in a completely analogous way. If $\psi = \ell\tau' \in \mathcal{O}$, then the inclusion $[1, \ell\tau'] \subseteq [1, \tau]$ is obvious. If $\psi = \ell\widehat{\tau'}$, then we can observe that the lattices $[1, \ell\tau']$ and $[1, \ell\widehat{\tau'}]$ coincide. Indeed, for example, we have

$$\tau' = \widehat{\tau'} \cdot (-1) + 1 \cdot 2T(\tau').$$

Hence, it is just a matter of choosing the generators of the lattice. Thus, we are left with the chain of inclusions

$$[1, \ell^2\tau] \subseteq [1, \ell\tau'] \subseteq [1, \tau].$$

This implies that

$$\ell^2 = [\mathcal{O} \colon [1, \ell^2\tau]] = [\mathcal{O} \colon [1, \ell\tau']][[1, \ell\tau'] \colon [1, \ell^2\tau]] = [\mathcal{O} \colon [1, \ell\tau']][\mathcal{O}' \colon [1, \ell\tau]].$$

Thus, we only have three possibilities.

If $[\mathcal{O} \colon [1, \ell\tau']] = \ell$, then

$$\ell = [\mathcal{O} \colon \mathcal{O}'][\mathcal{O}' \colon [1, \ell\tau']] = \ell[\mathcal{O} \colon \mathcal{O}'],$$

which implies $[\mathcal{O}\colon \mathcal{O}'] = 1$.

If $[\mathcal{O}\colon [1, \ell\tau']] = \ell^2$, then

$$\ell^2 = [\mathcal{O}\colon \mathcal{O}'][\mathcal{O}'\colon [1, \ell\tau']] = \ell[\mathcal{O}\colon \mathcal{O}'],$$

which implies $[\mathcal{O}\colon \mathcal{O}'] = \ell$.

If $[\mathcal{O}\colon [1, \ell\tau']] = 1$, then $[\mathcal{O}'\colon [1, \ell\tau]] = \ell^2$ and we can conclude $[\mathcal{O}'\colon \mathcal{O}] = \ell$, using the same argument of the previous case.

□

The next proposition is due to Kohel.

Proposition 3.1.6. *Let $\phi\colon E \longrightarrow E'$ be an isogeny of elliptic curves over K and let $\mathcal{O} = \mathrm{End}(E)$ and $\mathcal{O}' = \mathrm{End}(E')$. The orders $\mathcal{O}$ and $\mathcal{O}'$ are isomorphic if and only if there exists an isogeny $\psi\colon E \longrightarrow E'$ of degree coprime to* $\deg\phi$.

Proof. See [20], proposition 22. □

Exploiting this last result, we can prove the following proposition.

Proposition 3.1.7. *Let $\phi\colon E \longrightarrow E'$ be an isogeny between elliptic curves over K with complex multiplication by $\mathcal{O}$ and $\mathcal{O}'$, respectively. If $[\mathcal{O}\colon \mathcal{O}'] = \ell$, then ℓ divides the degree of ϕ.*

Proof. Suppose by contradiction that ℓ does not divide $\deg\phi$. The lattice inclusion $\mathcal{O}' \subseteq \mathcal{O}$ yields an ℓ-isogeny ψ between E and E'. Hence ψ is an isogeny of degree coprime to $\deg\phi$. By the previous theorem, this is equivalent to have $\mathcal{O} \simeq \mathcal{O}'$, which contradicts our hypothesis on the index of the inclusion of these orders. □

Theorem 3.1.8. *Let $E/\mathbb{C}$ be an elliptic curve with complex multiplication by $\mathcal{O}$, an order in a imaginary quadratic number field $L = \mathrm{End}^A(E)$ with discriminant d_L. Let ℓ be a prime number.*

(1) *If ℓ does not divide the conductor of $\mathcal{O}$, there are $1 + (\frac{d_L}{\ell})$ horizontal, $\ell - (\frac{d_L}{\ell})$ descending and no ascending ℓ-isogenies with domain E.*

(2) *If ℓ divides the conductor of $\mathcal{O}$, there are no horizontal, ℓ descending and one ascending ℓ-isogenies with domain E.*

Proof. Let f be the conductor of $\mathcal{O}$. Recall that ℓ-isogenies with domain E correspond to cyclic subgroups of order ℓ in $E[\ell]$, the ℓ-torsion subgroup of E, and we know exactly the structure of this torsion group, which is

$$E[\ell] = \frac{\mathbb{Z}}{\ell\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell\mathbb{Z}}.$$

There are precisely $\ell + 1$ cyclic subgroups of order ℓ in $E[\ell]$ and so there are precisely $\ell + 1$ ℓ-isogenies with domain E.

First, suppose that $E = E_\Lambda$, where $\Lambda = \mathcal{O}$ (or any other lattice homothetic to $\mathcal{O}$). All ℓ-isogenies we are looking for arise from some lattice inclusion $\Lambda' \subseteq \Lambda$ of index ℓ. Let

μ and τ be the same quantities as in Theorem 3.1.5, so that we have $\mathcal{O} = [1, \tau]$. For all such inclusions, we have that Λ' must be of the shape $\Lambda_i = [\ell, \tau + i]$, for $i = 0, ..., \ell - 1$, or $\Lambda_l = [1, \ell\tau]$, by Lemma 2.1.15. All horizontal isogenies are of the form

$$E_\Lambda \longrightarrow E_{\Lambda'},$$

where Λ' is a proper $\mathcal{O}$-ideal of index ℓ. By Proposition 2.3.15, if ℓ does not divide f, there are exactly $1 + (\frac{d_L}{\ell})$ such ideals. Moreover, in this case there is no order $\mathcal{O}'$ such that $[\mathcal{O}' \colon \mathcal{O}] = \ell$. Indeed, if such an order were to exist, we would have

$$f = [\mathcal{O}_L \colon \mathcal{O}'][\mathcal{O}' \colon \mathcal{O}] = \ell[\mathcal{O}_L \colon \mathcal{O}'],$$

which contradicts our hypothesis. Therefore, there is no ascending ℓ-isogeny and all the remaining $\ell - (\frac{d_L}{\ell})$ ℓ-isogenies must be descending.

On the other hand, if ℓ does divide f, there is no proper $\mathcal{O}$-ideal, so there is no horizontal ℓ-isogeny with domain E. Moreover, there exists an integer $e > 1$ such that $f = e\ell$, and we can define $\mathcal{O}' = \mathbb{Z} + e\mathcal{O}_L$, which is the unique order of index ℓ in $\mathcal{O}$. As a lattice, it can be represented by $[1, \tau']$, where $\tau' = e\mu$. Now, we have to check which of the Λ_i defined above are lattices such that $\mathrm{End}(E_{\Lambda_i}) = \mathcal{O}'$. The number of such lattices will correspond to the number of ascending ℓ-isogenies with domain E. First, notice that Λ_0 is fixed by $\mathcal{O}'$:

$$\tau' \cdot \ell = \tau \in \Lambda_0 \quad \text{and} \quad \tau' \cdot \tau = \ell(\tau')^2 = \ell(a\tau' - b) = a\tau - \ell b \in \Lambda_0.$$

We used that τ' satisfies the quadratic equation

$$(\tau')^2 - a\tau' + b = 0,$$

where $a = \mathrm{Tr}_{L/\mathbb{Q}}(\tau')$ and $b = \mathrm{N}_{L/\mathbb{Q}}(\tau')$. All the other Λ_i are not fixed by $\mathcal{O}'$:

$$\tau' \cdot (\tau + i) = \ell(\tau')^2 + i\tau' = a\ell\tau' - b\ell + i\tau' = (a\ell + i)\tau' - b\ell,$$

which is not in Λ_i for $i = 1, ..., \ell - 1$, and

$$1 \cdot \tau' = \tau',$$

which is not in Λ_l. Hence, if ℓ divides, f there are no horizontal, ℓ descending and one ascending ℓ-isogenies with domain E.

Suppose now that $E = E_\Lambda$, where Λ is (homothetic to) a proper $\mathcal{O}$-ideal $\mathfrak{a}$, which we can assume has prime norm $p \neq \ell$ by Proposition 2.3.16. The complex multiplication action of $\mathfrak{a}$ on $E = E_\mathfrak{a}$ gives an isogeny

$$\phi_\mathfrak{a} \colon E \longrightarrow E_{\mathfrak{a}^{-1}\mathfrak{a}} = E_\mathcal{O},$$

which is horizontal and whose degree is $\mathfrak{N}(\mathfrak{a}) = p$. Let $\phi \colon E \longrightarrow E'$ be an ℓ-isogeny and $\mathrm{End}(E') = \mathcal{O}'$ be an order in L. Let $\mathfrak{a}'$ be the ideal $\mathfrak{a}, \mathfrak{a} \cap \mathcal{O}'$ or $\mathfrak{a}\mathcal{O}'$, depending on whether ϕ is horizontal, descending or ascending, respectively. In each case, the ideal $\mathfrak{a}'$ is a proper $\mathcal{O}'$-ideal of norm p. Indeed, if the ϕ is horizontal, this is trivial. If ϕ is descending, by

Lemma 2.4.4, the $\mathcal{O}_L$-ideal $\mathfrak{a}\mathcal{O}_L$ has norm p and then $\mathfrak{a}\mathcal{O}_L \cap \mathcal{O}' = \mathfrak{a} \cap \mathcal{O}' = \mathfrak{a}'$ has norm p, again by Lemma 2.4.4. Using the same lemma, one can similarly work out the details in the case ϕ is ascending. Hence, $\mathfrak{a}'$ induces an horizontal p-isogeny

$$\phi_{\mathfrak{a}'}\colon E_{\mathfrak{a}'} \longrightarrow E'^{-1}_{\mathfrak{a}}a' = E_{\mathcal{O}'}.$$

We have the following diagram

$$\begin{tikzcd}
E \arrow[r, "\phi_{\mathfrak{a}}"] \arrow[d, "\phi"'] & E_{\mathcal{O}} \arrow[d, "\tilde{\phi}"] \\
E' \arrow[r, "\phi_{\mathfrak{a}'}"'] & E_{\mathcal{O}'}
\end{tikzcd}$$

where the isogeny $\tilde{\phi}\colon E_{\mathcal{O}} \longrightarrow E'_{\mathcal{O}}$ is the unique isogeny with kernel $\phi_{\mathfrak{a}}(\ker(\phi_{\mathfrak{a}'} \circ \phi))$, i.e. it is the isogeny that makes the diagram commute. Since both $\phi_{\mathfrak{a}}$ and $\phi_{\mathfrak{a}'}$ are horizontal, the ℓ-isogeny $\tilde{\phi}$ must be of the same type as ϕ. Then, the whole theorem follows from the special case we proved before. □

Thanks to the above results, we can observe that, given ℓ a prime number, isogeny graphs $G_{\mathbb{C},\ell}$ have a very rigid structure. Each component of the graph has infinite vertices and corresponds to some quadratic imaginary number field.

Definition 3.1.9. An *ℓ-volcano V* is an undirected graph whose vertices are divided into one or more levels $V_0, ..., V_d$ such that the following hold:

(1) the subgraph V_0, called the *surface*, is a regular graph of degree at most 2;

(2) for each $i > 0$, each vertex in V_i has exactly one neighbor lying in level V_{i-1} and no neighbors in level V_i;

(3) for each $i < d$, each vertex in V_i has degree $\ell + 1$.

The subgraph V_d is called the *floor* of the volcano, and the integer d is called *depth.* If the number of levels d is infinite, the third condition is empty and we say that V is an *infinite ℓ-volcano.*

An ℓ-volcano is completely determined by the three parameters ℓ, d and $|V_0|$. Our previous theorems and lemmas directly yield the following result.

Theorem 3.1.10. *Let ℓ be a prime number. Let V be any component of the isogeny graph $G_{\mathbb{C},\ell}$ not containing the j-invariants* 0 *and* 1728*. Then, if we consider V as an undirected graph, it is an infinite ℓ-volcano.*

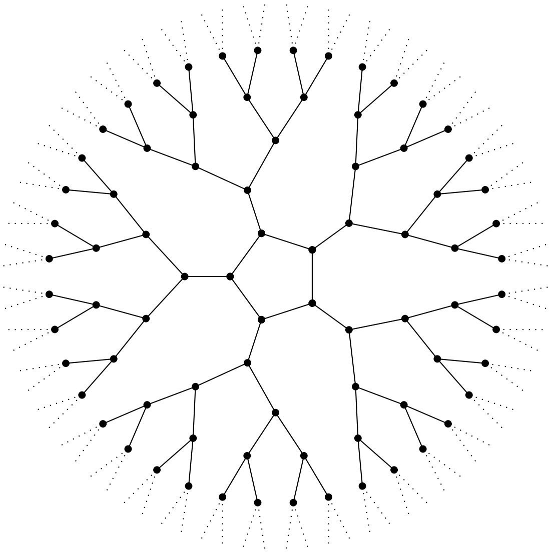

Figure 3.1: An infinite 2-volcano graph.

3.2 Ordinary Volcanoes

In this section, we will focus on finite field elliptic curves and prove that the ordinary components of isogeny graphs of fixed prime degree are volcanoes. Precisely, our strategy will be to transfer the previous results to finite fields. The next lemmas are the first steps in this direction.

Lemma 3.2.1. *Let $q = p^r$ and $t \not\equiv 0 \bmod p$ such that $|t| \leq 2\sqrt{q}$, for some $r > 0$ and p prime. Let $L = \mathbb{Q}[\sqrt{t^2 - 4q}]$ be an imaginary quadratic number field. Then there exist elliptic curves $E/\mathbb{F}_q$ such that $\operatorname{End}(E) = \mathbb{Z}[\pi] \subseteq L$, where π is the Frobenius endomorphism, defined as the root of the polynomial*

$$\pi^2 - t\pi + q = 0.$$

Let $\mathcal{O}$ be an order of discriminant d in L. Then $\operatorname{Ell}_{\mathcal{O}}(\mathbb{F}_q)$ is either empty or has full cardinality, which means that $\#\operatorname{Ell}_{\mathcal{O}}(\mathbb{F}_q) = h(\mathcal{O})$, where $h(\mathcal{O})$ is the class number of $\mathcal{O}$. Moreover, the set of elliptic curve over $\mathbb{F}_q$ with endomorphism ring $\mathcal{O}$ is non-empty if there exists v such that d satisfies the relation

$$4q = t^2 - v^2 d.$$

Proof. See [30, Theorem 4.3] for the first part. Notice that the condition $t \not\equiv 0 \mod p$ ensures that the curve E is an ordinary elliptic curve, by Theorem 1.6.2. From [23, Theorem 4.1], if $\mathcal{O}$ is an order in L such that $\operatorname{Ell}_{\mathcal{O}}(\mathbb{F}_q)$ is non-empty, we get that the Hilbert class polynomial splits completely over $\mathbb{F}_q$ and its roots are j-invariants of the elliptic curves over $\mathbb{F}_q$ with endomorphism ring $\mathcal{O}$. By Proposition 2.4.2, we have that $\operatorname{Ell}_{\mathcal{O}}(\mathbb{F}_q)$ has the desired cardinality. Next, we show why the last statement holds true. [30, Theorem 4.3] shows that all orders that contain the order $\mathbb{Z}[\pi]$ are endomorphism rings of a certain elliptic curve over $\mathbb{F}_q$. Recall that, in the proof of Corollary 1.6.9, we proved

$$d(\mathbb{Z}[\pi]) = t^2 - 4q.$$

Let $\mathcal{O}$ be an order in L. We have that $\mathcal{O}$ contains $\mathbb{Z}[\pi]$ if and only if its discriminant d satisfies the relation

$$d(\mathbb{Z}[\pi]) = [\mathcal{O}\colon \mathbb{Z}[\pi]]^2 d,$$

by Lemma 2.4.5. Thus, we have the relation

$$4q = t^2 - [\mathcal{O}\colon \mathbb{Z}[\pi]]^2 d$$

Hence, we have that if $\mathcal{O}$ is an imaginary order of discriminant d and there exists v such that the relation

$$4q = t^2 - v^2 d$$

holds, then the set of elliptic curves over $\mathbb{F}_q$ with endomorphism ring $\mathcal{O}$ is non-empty and has full cardinality. □

The next theorem is the finite field version of Theorem 3.1.8.

Theorem 3.2.2. *Let $E/\mathbb{F}_q$ be an elliptic curve with complex multiplication by $\mathcal{O}$, an order of discriminant d in the imaginary quadratic number field L with discriminant d_L. Assume d is coprime with q. Let ℓ be a prime natural number that does not divide q.*

(1) *If ℓ does not divide $[\mathcal{O}_L\colon \mathcal{O}]$, then E admits $1 + (\frac{d_L}{\ell})$ horizontal and zero ascending ℓ-isogenies over $\mathbb{F}_q$.*

(2) *If ℓ divides $[\mathcal{O}_L\colon \mathcal{O}]$, then E admits no horizontal and one ascending ℓ-isogeny over $\mathbb{F}_q$.*

In both cases, E can admit zero or $\ell - (\frac{d_L}{\ell})$ descending isogenies over $\mathbb{F}_q$, depending on whether $Ell_{\mathcal{O}'}(\mathbb{F}_q)$ is empty or has full cardinality, where $\mathcal{O}'$ is the order of index ℓ in $\mathcal{O}$.

Proof. Suppose $\phi\colon E \longrightarrow E'$ is a ℓ-isogeny over $\mathbb{C}$ with $\mathrm{End}(E) = \mathcal{O}$ and $\mathrm{End}(E') = \mathcal{O}'$. Suppose that $\mathrm{Ell}_{\mathcal{O}}(\mathbb{F}_q)$ and $\mathrm{Ell}_{\mathcal{O}'}(\mathbb{F}_q)$ are non-empty. We can look at ϕ as an isogeny over M, the biggest between ring class fields of $\mathcal{O}$ and $\mathcal{O}'$. Since $\mathcal{O} \subseteq \mathcal{O}'$ or $\mathcal{O}' \subseteq \mathcal{O}$ must hold, one ring class field must contain the other. This claim can be proved looking at how the correspondence between lattice inclusions and isogenies actually works, for example in the proof of [33, Theorem V.4.1.b]. Now, let $\mathfrak{q}$ be a prime ideal of norm q in $\mathcal{O}_M$. We can reduce the equations of E and E' through the map

$$\mathcal{O}_M \longrightarrow \mathcal{O}_M/\mathfrak{q} \simeq \mathbb{F}_q$$

and obtain a reduced isogeny $\overline{\phi}\colon \overline{E} \longrightarrow \overline{E'}$ of degree ℓ with $\mathrm{End}(\overline{E}) = \mathcal{O}$ and $\mathrm{End}(\overline{E'}) = \mathcal{O}'$, by Theorem 2.4.8. The degree of the reduced isogeny cannot change: ℓ does not divide q, so $E[\ell] \simeq \overline{E}[\ell]$, which implies $\ker\phi \simeq \ker\overline{\phi}$. Since ϕ is an isogeny over $\mathbb{C}$, it is separable and this implies that also $\overline{\phi}$ is separable. Conversely, if $\overline{\phi}\colon \overline{E} \longrightarrow \overline{E'}$ is a $\ell-$isogeny of curves over $\mathbb{F}_q$, by Theorem 2.4.9, we can lift these two curves to elliptic curves over L, preserving their endomorphism rings, and there is an ℓ-isogeny $\phi\colon E \longrightarrow E'$ that reduces to $\overline{\phi}$, whose kernel is isomorphic to $\ker\overline{\phi}$. Then, using Lemma 3.2.1, one immediately proves the claim about horizontal and ascending ℓ-isogenies. The claim about descending isogenies is true since their existence depends on whether $\mathrm{Ell}_{\mathcal{O}'}(\mathbb{F}_q)$ in non-empty or has full cardinality. □

We can now completely organize ℓ-isogenies from an elliptic curve $E/\mathbb{F}_q$ with complex multiplication by $\mathcal{O}$, which is an order in a quadratic imaginary number field L of discriminant d_L. The following table summarizes what we have proved so far.

First Condition	Second Condition	Horizontal	Ascending	Descending	Total
$\ell \nmid [\mathcal{O}_L\colon \mathcal{O}]$	$\ell \nmid [\mathcal{O}\colon \mathbb{Z}[\pi]]$	$1+(\frac{d_L}{\ell})$	0	0	$1+(\frac{d_L}{\ell})$
$\ell \nmid [\mathcal{O}_L\colon \mathcal{O}]$	$\ell \mid [\mathcal{O}\colon \mathbb{Z}[\pi]]$	$1+(\frac{d_L}{\ell})$	0	$\ell-(\frac{d_L}{\ell})$	$\ell+1$
$\ell \mid [\mathcal{O}_L\colon \mathcal{O}]$	$\ell \nmid [\mathcal{O}\colon \mathbb{Z}[\pi]]$	0	1	0	1
$\ell \mid [\mathcal{O}_L\colon \mathcal{O}]$	$\ell \mid [\mathcal{O}\colon \mathbb{Z}[\pi]]$	0	1	ℓ	$\ell+1$

We explain why the *Second Condition* is relevant for the number of descending ℓ-isogenies. If ℓ does not divide $[\mathcal{O}_L\colon \mathcal{O}]$ and $[\mathcal{O}\colon \mathbb{Z}[\pi]]$, then there is no order of index ℓ in $\mathcal{O}$. Indeed, assume by contradiction that $\mathcal{O}'$ is such an order. Then we would have

$$[\mathcal{O}\colon \mathbb{Z}[\pi]] = [\mathcal{O}\colon \mathcal{O}'][\mathcal{O}'\colon \mathbb{Z}[\pi]] = \ell[\mathcal{O}'\colon \mathbb{Z}[\pi]] \tag{3.1}$$

which gives a contradiction. If ℓ does not divide $[\mathcal{O}_L\colon \mathcal{O}]$ but divides $[\mathcal{O}\colon \mathbb{Z}[\pi]]$, then there exists an order of index ℓ in $\mathcal{O}$ and the precise number of descending isogenies is given by Theorem 3.2.2. The same argument can be applied to the third and fourth rows of the table. Observe that if ℓ^s divides exactly $[\mathcal{O}\colon \mathbb{Z}[\pi]]$ for some $s > 0$, then there is a chain of s ℓ-isogenies

$$\begin{tikzcd} E = E_0 \arrow[r, "\phi_1"] & E_1 \arrow[r, "\phi_2"] & \ldots \arrow[r, "\phi_s"] & E_s \end{tikzcd}$$

such that $\mathrm{End}(E_i) = \mathcal{O}_i$ is an order of index ℓ^i in $\mathcal{O} = \mathcal{O}_0$.

Remark 3.2.3. One could argue that the Frobenius endomorphism depends on the specific elliptic curve and that equality (3.1) is pointless, because on the left hand side we are considering π as an element of $\mathcal{O}$ and on the right hand side as an element of $\mathcal{O}'$. However, two elliptic curves are isogenous over K if and only if they have the same number of points over K. If K is the finite field with q elements, by Hasse estimate, two isogenous curves have the same t, which is the trace of the Frobenius endomorphism. Hence, after embedding their endomorphism rings in the same quadratic imaginary number field, their Frobenius endomorphism coincide, up to dualization. Indeed, the Frobenius endomorphism and its dual isogeny satisfy the equation

$$x^2 - tx + q = 0.$$

If we exclude components containing the two exceptional j-invariants 0 and 1728, the ordinary components of $G_{q,\ell}$ are all ℓ-volcanoes.

Theorem 3.2.4 (Kohel). *Let $\mathbb{F}_q$ be a finite field, ℓ a prime number not dividing q and V an ordinary component of $G_{q,\ell}$ that does not contain the j-invariants 0 or 1728. Let L be the endomorphism algebra shared by all curves in V. Then V is an ℓ-volcano of depth d with the following properties:*

(1) *for each $i \in \{0, ..., d\}$, the curves in level V_i all have the same endomorphism ring $\mathcal{O}_i$;*

(2) *the subgraph consisting of the surface V_0 has degree $1 + (\frac{d_0}{\ell})$, where d_0 is the discriminant of $\mathcal{O}_0$;*

(3) *if $(\frac{d_0}{\ell}) \geq 0$, then $|V_0|$ is the order of $[\mathfrak{l}]$ in $\mathcal{O}_0$, where $\mathfrak{l}$ is a prime ideal lying over ℓ. Otherwise, $|V_0| = 1$;*

(4) *the depth d is the positive integer that satisfies $4q = t^2 - \ell^{2d}v^2d_0$, where ℓ does not divide v, and t is the trace of the Frobenius endomorphism of the curves in V;*

(5) *the prime number ℓ does not divide the conductor $[\mathcal{O}_L \colon \mathcal{O}_0]$ and $[\mathcal{O}_i \colon \mathcal{O}_{i+1}] = \ell$ for each $0 \leq i < d$;*

Proof. The only automorphisms admitted by an ordinary elliptic curve E with j-invariant different from 0 and 1728 are $\pm 1 \in \mathrm{End}(E)$, by Theorem 1.5.13. Hence, every edge (j_1, j_2) appears with the same multiplicity as (j_2, j_1) and this allows us to consider V as an undirected graph. If the conductor of $\mathrm{End}(E) = \mathcal{O}$ is divisible by ℓ, then E admits an ascending ℓ-isogeny and so we can partition V into levels $V_0, ..., V_d$ with $j(E) \in V_i$ if and only if ℓ divides precisely i times the conductor of $\mathcal{O}$. The set of vertices V is finite, since going down there will be some endomorphism ring $\mathcal{O}'$ such that its conductor is divisible by ℓ but $[\mathcal{O}' \colon \mathbb{Z}[\pi]]$ is not. This implies that at some point there will be no descending ℓ-isogeny and d is bounded. This proves (1) and (5) and the theorems we proved before also prove (2) and the fact the V is an ℓ-volcano. If $(\frac{d_0}{\ell}) = -1$, then the subgraph V_0 has degree zero and so we must have $|V_0| = 1$. Otherwise, there exists a $\mathcal{O}_0$-proper ideal $\mathfrak{l}$ of norm ℓ and its ideal class acts on the vertices of V_0 via horizontal isogenies. This proves (3). If $4q = t^2 - \ell^{2d}v^2d_0$ with ℓ not dividing v, then the set $\mathrm{Ell}_{\mathcal{O}_i}$ must be non-empty for each $0 \leq i \leq d$ and the set $\mathrm{Ell}_{\mathcal{O}_{d+1}}$ must be empty since ℓ^{d+1} does not divide v. □

Using SageMath, we can compute and visualize isogeny graphs.

Example 3.2.5. Let E be the ordinary elliptic curve of j-invariant 32 over the field $\mathbb{F}_{113^2}$. On the left of Figure 3.2, we can see the 3-isogeny graph in which one of the vertices on the surface represents the elliptic curve E. If we denote by $\mathcal{O}_E$ the endomorphism ring of E, we can deduce that 3 does not divide the conductor of $\mathcal{O}_E$, since E is on the surface. On the other hand, if we denote the Frobenius endomorphism of E by π_E, we have that 3 divides precisely 2 times the index $[\mathcal{O}_E \colon \mathbb{Z}[\pi_E]]$. Moreover, we know from the third point of Theorem 3.2.4 that 3 is a split prime in the ring of integers and the order of any prime over 3 in the class group is exactly 8, the number of vertices on the surface of the volcano.

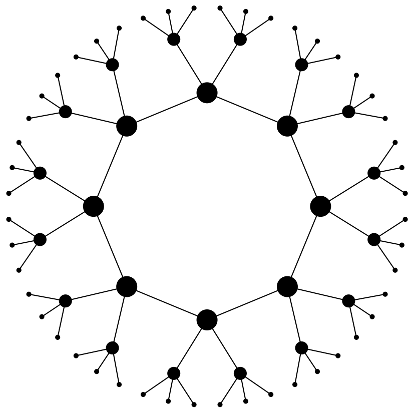

Figure 3.2: A 3-volcano arising from the ordinary elliptic curve of j-invariant 32 over the field $\mathbb{F}_{113^2}$.

Example 3.2.6. Let F be the ordinary elliptic curve of j-invariant 144 over the field $\mathbb{F}_{311^2}$. On the right of Figure 3.3, we can see the 5-isogeny graph in which one of the vertices on the surface represents the elliptic curve F. Let $\mathcal{O}_F$ and π_F be the endomorphism ring of F and the Frobenius endomorphism of F, respectively. Similarly to what we have said for E, we have that 5 does not divide the conductor of $\mathcal{O}_F$, but it divides precisely once the index $[\mathcal{O}_F\colon \mathbb{Z}[\pi_F]]$. Furthermore, 5 is a split prime in the ring of integers, and the order of each prime over 5 in the class group is exactly 13.

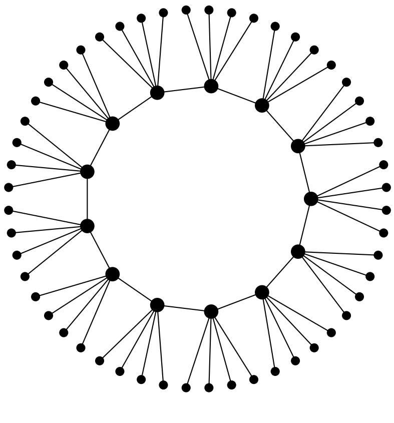

Figure 3.3: A 5-volcano arising from the ordinary elliptic curve of j-invariant 144 over the field $\mathbb{F}_{311^2}$.

Example 3.2.7. In Figure 3.4, we have a volcano with depth 0, thus a cycle. As in the previous cases, if we denote by $G/\mathbb{F}_{197}$ the curve of j-invariant 112, we have that 3 is a split prime in the ring of integers of the endomorphism algebra of G and the order of a prime over 3 in the class group is 6.

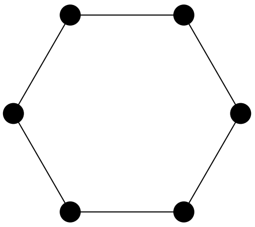

Figure 3.4: A 3-volcano arising from the ordinary elliptic curve of j-invariant 112 over the field $\mathbb{F}_{197}$. There is no descending isogeny, this means that the 3 does not divide the index of $\mathbb{Z}[\pi]$ in the endomorphism ring.

3.3 Supersingular Volcanoes

In general, isogeny graphs of supersingular elliptic curves over finite fields have an irregular structure: they are Ramanujan graphs. This kind of graphs are widely used in cryptography thanks to its rapid mixing properties. Recall that it is enough to consider j-invariants lying in $\mathbb{F}_{p^2}$ in order to cover all supersingular curves. However, if we restrict ourselves to j-invariants in $\mathbb{F}_p$, isogenies defined over $\mathbb{F}_p$ and isomorphisms defined over $\mathbb{F}_p$, then the supersingular isogeny graph becomes considerably smaller, and we will prove that it has a rigid structure similar to the one in the ordinary case.

Lemma 3.3.1. *Let $p > 3$ be a prime number, and let S_p be the set of all supersingular j-invariants in $\mathbb{F}_p$. Then we have*

$$\#S_p = \begin{cases} \frac{1}{2}h(-4p) & \text{if } p \equiv 1 \bmod 4, \\ h(-p) & \text{if } p \equiv 7 \bmod 8, \\ 2h(-p) & \text{if } p \equiv 3 \bmod 8, \end{cases}$$

where $h(d)$ is the class number of the imaginary quadratic number field $\mathbb{Q}(\sqrt{d})$.

Proof. For a slightly more general proof, see [10, Theorem 14.18]. □

Given $E/\mathbb{F}_p$ an elliptic curve, we will denote simply by $\mathrm{End}_p(E)$ the $\mathbb{F}_p$-rational endomorphism ring of E. The following theorem represents a key point in the study of supersingular components of $G_{p,\ell}$.

Theorem 3.3.2. *Let $p > 3$ be a prime number and E be a supersingular elliptic curve over $\mathbb{F}_p$. Then, denoting by π the Frobenius endomorphism of E, we have $\mathrm{End}^A(E) = \mathbb{Q}(\pi)$ and $\mathrm{End}_p(E)$ is an order in $\mathbb{Q}(\pi)$ with conductor prime to p.*

Proof. The proof can be found in [14] or in [37]. □

Recall that the number of $\mathbb{F}_p$-rational points of any supersingular curve is $p+1$ and so the trace of the Frobenius endomorphism is zero. Hence, the relation $\pi^2 + p = 0$ holds, and we have that $L = \mathbb{Q}(\pi) = \mathbb{Q}(\sqrt{-p})$. Moreover, we have the chain of inclusions

$$\mathbb{Z}[\pi] = \mathbb{Z}[\sqrt{-p}] \subseteq \mathrm{End}_p(E) = \mathcal{O} \subseteq \mathcal{O}_L = \mathbb{Z}\left[\frac{d_L + \sqrt{d_L}}{2}\right],$$

where d_L is the discriminant of L. If d is the discriminant of $\mathcal{O}$, then we have $d = f^2 d_L$, where f is the conductor of $\mathcal{O}$. Using basic number theory, if $p \equiv 1 \mod 4$, then we have $d = d_L = -4p$, so that $\mathbb{Z}[\pi] = \mathcal{O}_L$ and $\mathcal{O}$ is forced to be $\mathcal{O}_L$. On the other hand, if $p \equiv 3 \mod 4$ then we have $d_L = -p$, so that $\mathbb{Z}[\pi]$ has conductor 2 in $\mathcal{O}_L = \mathbb{Z}\left[\frac{1+\sqrt{-p}}{2}\right]$ and $\mathcal{O}$ must be one of those two orders. In terms of isogeny volcanoes, we can say that $G_{p,\ell}$ admits only two levels.

By Theorem 3.1.7, since $[\mathrm{End}_p(E)\colon \mathrm{End}_p(E')] \in \{1,2\}$, we have that any non-horizontal isogeny between supersingular elliptic curves over $\mathbb{F}_p$ has degree divisible by 2.

Let $p > 3$. A supersingular elliptic curve over $\mathbb{F}_p$ has $p+1$ points and so all quadratic twists have the same number of points. Thus, twists are isogenous but lie in different $\mathbb{F}_p$-isomorphism classes. Moreover, a priori, we should also deal with cubic, quartic and sextic twits for j-invariants 0 and 1728. The next lemma proves that this is not the case.

Lemma 3.3.3. *Suppose $p > 3$ and $j \in \mathbb{F}_p$. Let $C_{p,j}$ be the set of elliptic curves over $\mathbb{F}_p$ with j-invariant j up to $\mathbb{F}_p$-isomorphism. Then we have*

$$\#C_{p,j} = \begin{cases} 6 & \textit{if } j = 0 \textit{ and } p \equiv 1 \mod 3, \\ 4 & \textit{if } j = 1728 \textit{ and } p \equiv 1 \mod 4 \\ 2 & \textit{otherwise.} \end{cases}$$

Proof. This follows from Theorem 1.5.13. For a precise proof, see [3, Theorem 2.2]. □

The curves of j-invariant 0 and 1728 are supersingular if and only if $p \equiv 2 \mod 3$ and $p \equiv 3 \mod 4$, respectively. See [33, Examples V.4.4, V.4.5]. Thus, given a j-invariant, there are exactly two $\mathbb{F}_p$-isomorphism classes of elliptic curves over $\mathbb{F}_p$ with that j-invariant.

Proposition 3.3.4. *Let $p > 3$ and let E be a supersingular elliptic curve over $\overline{\mathbb{F}_p}$. Then E is defined over $\mathbb{F}_p$ if and only if $\mathbb{Z}[\sqrt{-p}] \subseteq \mathrm{End}(E)$.*

Proof. Let $\phi \in \mathrm{End}(E)$ satisfy $\phi^2 = -p$. Then ϕ is an isogeny of degree p and $\hat{\phi} \circ \phi = p$. Since E is supersingular, it follows that ϕ has kernel O and so is purely inseparable. Therefore, by Lemma 1.2.19, there exists a separable isogeny $\psi\colon E^{(p)} \longrightarrow E$ such that $\phi = \psi \circ \pi_p$, where π_p is the p-Frobenius morphism. We must have $\deg(\psi) = 1$, so that E and $E^{(p)}$ are isomorphic. This implies that $j(E) = j(E^{(p)}) = j(E)^p$ and thus E is defined over $\mathbb{F}_p$. The vice versa is trivial, since $\pi_p = \sqrt{-p}$ lies in $\mathrm{End}(E)$. □

As in the ordinary case, there is an intimate connection between supersingular elliptic curves over prime fields and certain elliptic curves in characteristic 0.

Proposition 3.3.5. *There is a one to one correspondence supersingular elliptic curves over $\mathbb{F}_p$ and elliptic curves E over the complex field with* $\mathrm{End}(E) \in \{\mathbb{Z}[\sqrt{-p}], \mathcal{O}_L\}$.

Proof. For the complete proof, see [13, Proposition 2.5]. Let $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ and $\mathrm{Super}(\mathbb{F}_p)$ be the set of all elliptic curves over $\mathbb{C}$ with complex multiplication over $\mathcal{O} \in \{\mathbb{Z}[\sqrt{-p}], \mathcal{O}_L\}$

and the set of $\mathbb{F}_p$-isomorphism classes of supersingular elliptic curves over $\mathbb{F}_p$, respectively. The idea of the proof is to show that there exists a bijective map

$$\operatorname{Ell}_{\mathcal{O}} \mathbb{C} \longrightarrow \operatorname{Super}(\mathbb{F}_p), \quad [E] \longmapsto [\overline{E}]$$

where $\overline{E}$ is the reduction of E modulo a fixed prime ideal over p. Surjectvity comes from Theorem 2.4.9 and injectivity comes from a counting argument based on Lemmas 3.3.3 and 3.3.1. □

Isogenies behave well under reduction. From [32, Theorem II.4.4], we know that reduction of isogenies is injective and preserves the degree. The next proposition claims that the reduction preserves also the $\mathbb{F}_p$-rationality.

Proposition 3.3.6. *Let $\overline{E_1}, \overline{E_2}$ be two supersingular elliptic curves over $\mathbb{F}_p$ and let (E_1, ψ), (E_2, ψ) be the Deuring lifts of $(\overline{E_1}, \pi)$ and $(\overline{E_2}, \pi)$, respectively. If there is an isogeny $\phi\colon E_1 \longrightarrow E_2$, then the reduced isogeny $\overline{\phi}\colon \overline{E_1} \longrightarrow \overline{E_2}$ is defined over $\mathbb{F}_p$.*

Proof. See [13, Proposition 2.6]. □

At this point, we have all the ingredients to prove that the supersingular isogeny graphs over prime fields are volcanoes.

Let $\overline{E}$ be a supersingular elliptic curve over $\mathbb{F}_p$. By Theorem 3.3.2, its endomorphism ring over $\mathbb{F}_p$ is an order $\mathcal{O} \in \{\mathbb{Z}[\sqrt{-p}], \mathcal{O}_L\}$ of discriminant $d \in \{-4p, -p\}$ in the number field $L = \mathbb{Q}(\sqrt{-p})$. Using Deuring lifting theorem, $\overline{E}$ can be lifted to an elliptic curve E over some number field with endomorphism ring $\operatorname{End}(E) = \mathcal{O}$. Let ℓ be a prime natural number. Let f be the conductor of $\mathcal{O}$ in $\mathcal{O}_L$. By Theorems 3.1.5 and 3.1.8, if ℓ divides f, then there are one ascending and ℓ descending ℓ-isogenies. If ℓ does not divide f, then we have $1 + (\frac{d_L}{\ell})$ horizontal and $\ell - (\frac{d_L}{\ell})$ descending ℓ-isogenies. In our case, as $\mathcal{O} \in \{\mathbb{Z}[\sqrt{-p}], \mathcal{O}_L\}$, the only possibilities are $f = 1$ and, if $p \equiv 3 \mod 4$, $f = 2$. In this second case, by Theorem 3.1.7, the only prime producing ascending or descending isogenies is $\ell = 2$.

We can build an infinite volcano of elliptic curves whose endomorphism ring is an order in L. If we consider the reduction modulo a prime ideal of norm p in $\mathcal{O}_L$, then all reduced curves are supersingular, because they all have trace zero. However, only a small part of the volcano survives. Indeed, by Proposition 3.3.4, all curves with endomorphism ring strictly contained in $\mathbb{Z}[\sqrt{-p}]$ do not reduce to elliptic curves over $\mathbb{F}_p$. For example, for $m \in \mathbb{N}$, the elliptic curves corresponding to the lattices $\mathbb{Z}[m\sqrt{-p}]$ have endomorphism ring $\mathbb{Z}[m\sqrt{-p}]$, which is strictly contained than $\mathbb{Z}[\sqrt{-p}]$. So these curves will not reduce to supersingular curves.

Now, we perform the transition from $G_{\mathbb{C},\ell}$ to $G_{p,\ell}$. By Theorem 3.3.6, the isogenies of the graph $G_{\mathbb{C},\ell}$ between curves with good reduction reduce to $\mathbb{F}_p$-rational outgoing isogenies. The only missing piece in the picture is showing that every isogeny between supersingular elliptic curves over $\mathbb{F}_p$ can be obtained via this reduction.

Lemma 3.3.7. *Let E be a supersingular elliptic curve over $\mathbb{F}_p$. Let $\ell > 2$ be a prime natural number. The number of $\mathbb{F}_p$-rational ℓ-isogenies with domain E is two if $(\frac{-p}{\ell}) = 1$ and zero otherwise. The number of $\mathbb{F}_p$-rational 2-isogenies with domain E is one if $p \equiv 1 \mod 4$ and three if $p \equiv 3 \mod 4$.*

Proof. Let E be a supersingular elliptic curve over $\mathbb{F}_p$. The $\mathbb{F}_p$-rational ℓ-isogenies correspond to Galois-invariant cyclic subgroups of $E[\ell]$. Indeed, each of these ℓ-isogenies is uniquely determined by its kernel, which must be a cyclic subgroup of $E[\ell]$. The condition of Galois-invariance ensures we are considering only $\mathbb{F}_p$-rational isogenies, by Remark 1.2.23. Let $\ell \neq p$, so that

$$E[\ell] = \frac{\mathbb{Z}}{\ell\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell\mathbb{Z}}$$

is a 2-dimensional vector space over $\mathbb{F}_\ell$. The Frobenius endomorphism π acts linearly on the ℓ-torsion subgroup. Fixing a basis for $E[\ell]$, the action of π is represented by a 2×2 matrix, whose characteristic polynomial modulo ℓ is $\pi^2 + p = 0$. This polynomial can be irreducible, split into two distinct linear factors, or split as a square of a linear polynomial.

Assume there is a cyclic subgroup G of $E[\ell]$ generated by a point P which is also Galois-invariant, i.e. $\pi(G) = G$. It follows that $\pi(P) = [a]P$, for some integer a. Thus, the linear map π has an eigenspace with eigenvalue a and the characteristic polynomial splits with linear factor $\pi - a$. Using the point P and completing it to a basis of $E[\ell]$, the only two possibilities are

$$\begin{pmatrix} a & b \\ 0 & a \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix},$$

for some integer $b \in \mathbb{N}$. The number of Galois-invariant cyclic subgroups of $E[\ell]$ in the left case is $\ell + 1$ or 1, depending on whether the matrix's upper left entry is zero or not. In the case on the right, we have two of them and if the polynomial is irreducible there are none. The polynomial $x^2 + p \bmod \ell$ can only have a repeated root for $\ell = 2$. When b is congruent to 0 (resp. 1) modulo 2, we have 3 (resp. 1) Galois-invariant subgroups of $E[2]$. We want to show that each possibility corresponds to the desired case, i.e. that when $\operatorname{End}_p E = \mathbb{Z}[\frac{-p+\pi}{2}]$ we have $b \equiv 0$ modulo 2, and when $\operatorname{End}_p E = \mathbb{Z}[\pi]$ we have $b \equiv 1$ modulo 2.

We have $b \equiv 0$ modulo 2 if and only if $\pi(P) = P$ and $\pi(Q) = Q$, so that $E[2] = \ker[2]$ is included in $\ker(\pi - 1)$. Since the multiplication by 2 map is separable, there exists a unique isogeny $\phi \in \operatorname{End}(E)$ such that $\pi - 1 = 2\phi$, see [33, Corollary III.4.1]. Moreover, ϕ is $\mathbb{F}_p$-rational since it is the quotient of two $\mathbb{F}_p$-rational maps, namely $\pi - 1$ and the multiplication by 2.

For any other ℓ, we get no $\mathbb{F}_p$-rational isogeny when the polynomial is irreducible and two when it splits, i.e. when $(\frac{-p}{\ell}) = 1$. □

Eventually, we proved that the numbers of ℓ-isogenies in characteristic 0 equals the number of $\mathbb{F}_p$-rational ℓ-isogenies. Since isogenies in characteristic 0 reduce to $\mathbb{F}_p$-rational isogenies, there is a one-to-one correspondence betweeen $\mathbb{F}_p$-rational ℓ-isogenies of supersingular elliptic curves over $\mathbb{F}_p$ and ℓ-isogenies of elliptic curves over $\mathbb{C}$ with endomorphisms ring containing $\mathbb{Z}[\sqrt{-p}]$.

We can sum up the above lines in the following theorem.

Theorem 3.3.8. *Let $p > 3$.*

(1) *If $p \equiv 1 \mod 4$, there are $h(-4p)$ $\mathbb{F}_p$-isomorphism classes of supersingular elliptic curves, all having the same endomorphism ring $\mathbb{Z}[\sqrt{-p}]$. From each of them, there*

are one horizontal 2-isogeny and two horizontal ℓ-isogenies for every prime $\ell > 2$ with $(\frac{-p}{\ell}) = 1$, all of them $\mathbb{F}_p$-rational.

(2) *If $p \equiv 3 \mod 4$, there are two levels in the supersingular isogeny graph. From each vertex, there are two horizontal ℓ-isogenies for every prime $\ell > 2$ with $(\frac{-p}{\ell}) = 1$.*

The second case splits into two subcases.

(2*a*) *If $p \equiv 7 \mod 8$, there are $h(-p)$ vertices on each level. We have a 2-isogeny from each vertex on the surface going down to the floor and two horizontal 2-isogenies.*

(2*b*) *If $p \equiv 3 \mod 8$, we have $h(-p)$ vertices on the surface and $3h(-p)$ on the floor. Each vertex on the surface have three 2-isogenies going down to the floor and no horizontal 2-isogeny.*

Example 3.3.9. Let $p = 71$. Over the field $\mathbb{F}_{71^2}$, we expect to find $\lfloor \frac{71}{12} \rfloor + 2 = 7$ j-invariants of supersingular elliptic curves, by Corollary 1.6.6. Recall that the j-invariant is a good representative of the isomorphism class over the algebraic closure of the base field. If we restrict ourselves to prime fields, there are more than one $\mathbb{F}_p$-isomorphism class for each j-invariant. Indeed, by Lemma 3.3.3, we expect that the number of $\mathbb{F}_{71}$-isomorphism classes is exactly 14.

Since $71 \equiv 7 \mod 8$, we know that the 2-isogeny graphs over $\mathbb{F}_{71}$ has two levels, with 7 edges on each level. Each vertex in the surface has three edges that correspond to two horizontal isogenies and one descending isogeny. In Figure 3.5, we have on the left the supersingular 2-isogeny graph over $\mathbb{F}_{71^2}$ and on the right the supersingular 2-isogeny graph over $\mathbb{F}_{71}$, with its much more regular structure.

Since $71 \equiv 3 \mod 4$ and $(\frac{-71}{3}) = 1$, we have that each vertex in the 3-isogeny graph over $\mathbb{F}_{71}$ has two horizontal isogenies.We still have two levels in the graphs, but they are disconnected. In Figure 3.6, on the left we have the $\mathbb{F}_{71^2}$-rational supersingular 3-isogeny graph, on the right the $\mathbb{F}_{71}$-rational supersingular 3-isogeny graph.

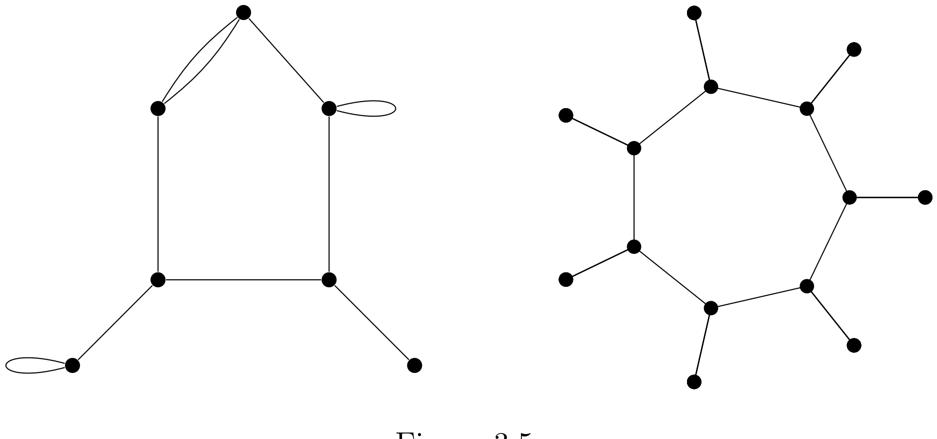

Figure 3.5

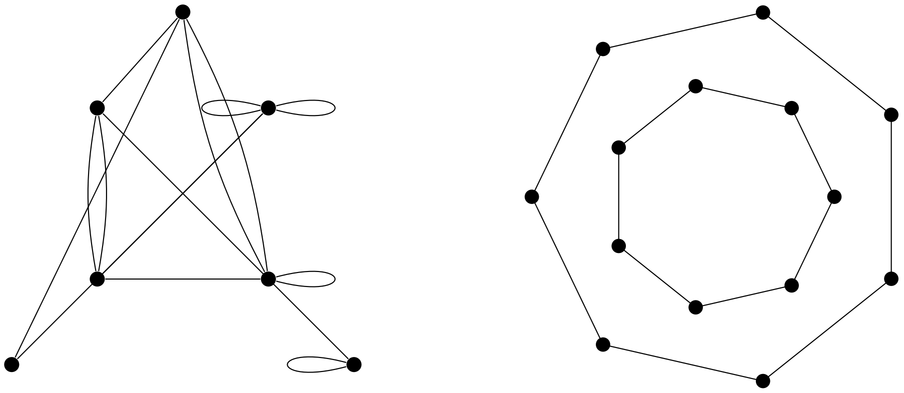

Figure 3.6

3.4 Expander Isogeny Graphs

In the previous sections, we focused on understanding the structure of isogeny graphs where the allowed degree for the isogenies is a single fixed prime number. In the following section, we want to understand the behavior of isogeny graphs when we consider isogenies of prime degree up to a given bound. We will show that, under the correct assumptions, these graphs become expander graphs.

First, we briefly recall some ideas from graph theory.

Definition 3.4.1. Let $G = (\mathcal{V}, \mathcal{E})$ be an undirected graph, where $\mathcal{V} = \{v_1, ..., v_n\}$ is the set of its vertices and $\mathcal{E}$ is the set of its edges. The *adjacency matrix* A of G is the $n \times n$ matrix such that the (i, j)-th entry is 1 if there exists an edge connecting v_i and v_j and 0 otherwise.

Recall that an undirected graph G is said to be k-regular if each vertex has degree exactly k. Since the graph G is undirected, the adjacency matrix A is symmetric, and this implies that it has n real eigenvalues that we will denote by

$$\lambda_1 \geq \cdots \geq \lambda_n.$$

We will often refer to the $\lambda_1, ..., \lambda_n$ as the eigenvalues of G rather than the eigenvalues of A, just to ease the notation. Notice that a function on the vertices of G

$$f \colon \mathcal{V} \longrightarrow \mathbb{C}$$

can be identified with a vector in $\mathbb{C}^n$. Let $\ell^2(\mathcal{V})$ be the finite dimensional complex Hilbert space of square-summable functions

$$f \colon \mathcal{V} \longrightarrow \mathbb{C}$$

with norm

$$\|f\| = \left(\sum_{v \in \mathcal{V}} |f(v)|^2 \right)^{1/2}$$

and inner product

$$\langle f, g\rangle = \sum_{v\in\mathcal{V}} f(v)\overline{g(v)}.$$

We can think of A as a self-adjoint operator in $\ell^2(\mathcal{V})$

$$\begin{aligned} A\colon \ell^2(\mathcal{V}) &\longrightarrow \ell^2(\mathcal{V}) \\ f &\longmapsto Af. \end{aligned}$$

Hence, we have

$$Af(v) = \sum_{w\in\mathcal{V}\colon (v,w)\in\mathcal{E}} f(w).$$

Lemma 3.4.2. *Let $G = (\mathcal{V}, \mathcal{E})$ be a k-regular graph. Then its largest and smallest eigenvalues λ_1 and λ_n satisfy*

$$k = \lambda_1 \geq \lambda_n \geq k.$$

Proof. One can immediately verify that k is always an eigenvalue of G, since it corresponds to the eigenvector $\mathbf{1} = (1, ..., 1)$. Moreover, the operator norm of A is exactly k, so all its eigenvalues need to have norm smaller than k. Indeed, if $f, g \in \ell^2(\mathcal{V})$ have both unitary norm, then

$$\begin{aligned} |\langle Af, g\rangle| =& \left| \sum_{v,w\in\mathcal{V}\colon (v,w)\in\mathcal{E}} f(v)\overline{g(w)} \right| \\ \leq& \frac{1}{2} \sum_{v,w\in\mathcal{V}\colon (v,w)\in\mathcal{E}} |f(v)|^2 + |g(w)|^2 \\ \leq& \frac{1}{2} k\|f\| + \frac{1}{2} k\|g\| = k. \end{aligned}$$

The first inequality comes from the relation

$$\begin{aligned} \|f - g\| &\geq 0 \\ \langle f - g, f - g\rangle &\geq 0 \\ \|f\|^2 + \|g\|^2 &\geq 2\langle f, g\rangle. \end{aligned}$$

□

Definition 3.4.3. Let $\varepsilon > 0$ and $k \geq 1$. A k-regular graph is called (one-sided) *ε-expander* if

$$\lambda_2 \leq (1 - \varepsilon)k.$$

and a *two-sided ϵ-expander* if it also satisfies

$$\lambda_n \geq -(1 - \varepsilon)k.$$

Most of the time, we will simply say that a certain graph is an expander graph, without specifying for which ε the above condition holds.

Recall that, given any $S \subseteq \mathcal{V}$, the characteristic function χ_S of S is the function such that $\chi_S(v) = 1$ if $v \in S$, $\chi_S(v) = 0$ else.

Lemma 3.4.4. *Let G be an undirected graph whose set of vertices is $\mathcal{V}$/ Given $S, T \subseteq \mathcal{V}$ and the adjacency matrix A of the graph G, the number of paths of length $t \geq 0$ in G starting from a vertex in T ending in a vertex of S is given by the inner product*

$$\langle \chi_S, A^t \chi_T \rangle.$$

Proof. In order to prove it, we can assume without loss of generality that $T = \{v\}$, for some $v \in \mathcal{V}$. Indeed, by linearity we have $A^t\chi_T = \sum_{x \in T} A^t \chi_{\{x\}}$, and to recover the number of paths of length t starting from T ending in S we can sum the numbers of paths of length t starting from each vertex of T ending in S. Analogously, we can assume without loss of generality that $S = \{v'\}$. We will proceed by induction. For $t = 0$, notice that $\langle \chi_{\{v'\}}, \chi_{\{v\}} \rangle$ equals one if $v = v'$, zero otherwise. We can look at it as the number of paths from v to v' of length zero. For $t = 1$, notice that $A\chi_{\{v\}}(x)$ equals one if $x \in \mathcal{V}$ and v are connected by an edge, zero else. Hence, $\langle \chi_{\{v'\}}, A\chi_{\{v\}} \rangle = \langle \chi_{\{v'\}}, \chi_U \rangle$ equals one if $v' \in U$, zero otherwise, where U is the set of vertices that are reachable from v with a path of length 1. This coincides with the number of possible paths of length 1 starting from v ending in v'. Let us now assume that $\langle \chi_{\{v'\}}, A^l \chi_{\{v\}} \rangle$ counts the number of paths of length $l \leq t - 1$ starting from v ending in v'. Since G is an undirected graph, A is symmetric and we have that

$$\langle \chi_{\{v'\}}, A^t \chi_{\{v\}} \rangle = \langle \chi_{\{v'\}} A, A^{t-1} \chi_{\{v\}} \rangle = \langle \chi_R, A^{t-1} \chi_{\{v\}} \rangle = \sum_{x \in R} \langle \chi_{\{x\}}, A^{t-1} \chi_{\{v\}} \rangle,$$

by linearity, where R is the set of vertices that are connected to v' by one edge. By inductive hypothesis, we have that, for each $x \in R$, $\langle \chi_{\{x\}}, A^{t-1} \chi_{\{v\}} \rangle$ is the number of possible paths of length $t-1$ starting from v ending in x. The sum of those numbers gives the number of possible paths of length $t - 1$ starting from v ending in R, which equals to the number of paths of length t starting from v ending in v', as we wanted to show. □

Expander graphs are a fundamental tool used in many areas of computer science thanks to their *rapid mixing property*, which we will prove in the following theorem. In the following, we will denote by log the natural logarithm.

Theorem 3.4.5. *Let G be a finite k-regular two-sided expander graph, for which the nontrivial eigenvalues are bounded by $|\lambda_i| < c$ for some constant $c < k$. Let S be any subset of $\mathcal{V}$ and let $v \in \mathcal{V}$. Then, a random walk of length at least*

$$\frac{\log \frac{2\#\mathcal{V}}{\#S^{1/2}}}{\log \frac{k}{c}}$$

starting from v will end in S with probability between $\frac{\#S}{2\#\mathcal{V}}$ and $\frac{3\#S}{2\#\mathcal{V}}$.

Proof. Let χ_S and $\chi_{\{v\}}$ be respectively the characteristic functions of the sets S and $\{v\}$ and let A be the adjacency matrix of G. The number of paths of length t that start in v and end in S is given by the inner product $\langle \chi_S, A^t\chi_{\{v\}}\rangle$. Let C be the subspace of $\ell^2(\mathcal{V})$ consisting of all constant functions. Its orthogonal complement is the subspace

$$C^\perp = \{f \in \ell^2(\mathcal{V})\colon \sum_{v\in\mathcal{V}} f(v) = 0\},$$

and the operator A preserves this subspace. Indeed, if $f \in C^\perp$, we have

$$Af = \left(\sum_{w\in\mathcal{V}} A_{v,w} f(w)\right)_{v\in\mathcal{V}}$$

and we can verify that it is still in the subspace $C^\perp$

$$\sum_{v\in\mathcal{V}}\sum_{w\in\mathcal{V}} A_{v,w} f(w) = \sum_{w\in\mathcal{V}}\sum_{v\in\mathcal{V}} A_{v,w} f(w) = k\sum_{w\in\mathcal{V}} f(w) = 0.$$

Moreover, by assumption, the norm operator of A is bounded by c on $C^\perp$. Let P be the projection from $\ell^2(\mathcal{V})$ onto $C^\perp$ and Q be the projection from $\ell^2(\mathcal{V})$ onto C. Then we have

$$Q\chi_S = \langle \frac{\mathbf{1}}{\sqrt{\#\mathcal{V}}}, \chi_S\rangle \frac{\mathbf{1}}{\sqrt{\#\mathcal{V}}} = \left(\sum_{v\in S}\frac{1}{\sqrt{\#\mathcal{V}}}\right)\frac{\mathbf{1}}{\sqrt{\#\mathcal{V}}} = \frac{\#S}{\#\mathcal{V}}\mathbf{1},$$

and similarly

$$Q\chi_{\{v\}} = \frac{\mathbf{1}}{\#\mathcal{V}}.$$

Thus, we can compute the number of desired possible paths

$$\begin{aligned}\langle \chi_S, A^t\chi_{\{v\}}\rangle =& \langle P\chi_S + Q\chi_S, A^t(P\chi_{\{v\}} + Q\chi_{\{v\}})\rangle \\ =& \langle P\chi_S, A^tP\chi_{\{v\}})\rangle + \left\langle \frac{\#S}{\#\mathcal{V}}\mathbf{1}, A^t\frac{\mathbf{1}}{\#\mathcal{V}}\right\rangle \\ =& \langle P\chi_S, A^tP\chi_{\{v\}})\rangle + \sum_{v\in\mathcal{V}}\frac{\#S}{\#\mathcal{V}}\frac{k^t}{\#\mathcal{V}} \\ =& \langle P\chi_S, A^tP\chi_{\{v\}})\rangle + \frac{\#S}{\#\mathcal{V}}k^t.\end{aligned}$$

Moreover, as one may expect, the number of all possible paths of length t starting from v is

$$\langle \mathbf{1}, A^t\chi_{\{v\}}\rangle = k^t.$$

The probability we aim to estimate is then

$$\mathbb{P} = \frac{\#S}{\#\mathcal{V}} + \frac{\langle P\chi_S, A^tP\chi_{\{v\}})\rangle}{k^t}.$$

Notice that the numerator of the latter term is bounded by

$$|\langle P\chi_S, A^t P\chi_{\{v\}}\rangle| \leq c^t \|P\chi_S\| \|P\chi_{\{v\}}\| \leq c^t \|\chi_S\| \|\chi_{\{v\}}\| = c^t \# S^{1/2},$$

where the last inequality holds because the projection is a continuous linear operator. Hence, we have that

$$\frac{\#S}{\#\mathcal{V}} - \left(\frac{c}{k}\right)^t \#S^{1/2} \leq \mathbb{P} \leq \frac{\#S}{\#\mathcal{V}} + \left(\frac{c}{k}\right)^t \#S^{1/2}. \tag{3.2}$$

Now, we want to compute the length $\bar{t}$ such that

$$\left(\frac{c}{k}\right)^{\bar{t}} \#S^{1/2} = \frac{\#S}{2\#\mathcal{V}}.$$

Simple computations lead to

$$\bar{t} = \frac{\log \frac{2\#\mathcal{V}}{\#S^{1/2}}}{\log \frac{k}{c}}.$$

Thus, if we perform a walk in the graph of $t > \bar{t}$, then we have

$$\left(\frac{c}{k}\right)^{t} \#S^{1/2} \leq \frac{\#S}{2\#\mathcal{V}}.$$

Exploiting this new inequality, from (3.2) we get

$$\frac{\#S}{2\#\mathcal{V}} \leq \mathbb{P} \leq \frac{3\#S}{2\#\mathcal{V}},$$

that is exactly the relation we were looking for. □

Next, we want to introduce a special kind of graphs that can be produced from a group.

Definition 3.4.6. Let G be a group generated by a subset S that is closed under inversion. We define its corresponding *Cayley graph* $Cay(G, S)$, whose vertices are the elements of G and $g, h \in G$ are connected by an edge if and only if there exists $s \in S$ such that $h = sg$.

Definition 3.4.7. Let G be a group. A *(complex multiplicative) character for* G is a group morphism

$$\chi \colon G \longrightarrow \mathbb{C}^*.$$

Let A be the adjacency matrix of the Cayley graph $Cay(G, S)$. If G is a finite abelian group, the eigenfunctions of A are precisely the characters $\chi \colon G \longrightarrow \mathbb{C}^*$. Indeed, for each $g \in G$ we have

$$A\chi(g) = \sum_{s\in S} \chi(sg) = \lambda_\chi \chi(g), \quad \text{where} \quad \lambda_\chi = \sum_{s \in S} \chi(s).$$

Hence, the spectrum of the adjacency matrix consists of character sums over the generating set S and the trivial eigenvalue comes from the trivial character $\mathbf{1}$.

Definition 3.4.8. Let $\mathcal{O}$ be an order in the quadratic imaginary number field L. Let χ be a character of $\mathrm{cl}(\mathcal{O})$. For all $s \in \mathbb{C}$, we define the formal *Hecke $\mathcal{L}$-function* of the character χ to be

$$\mathcal{L}_{\mathcal{O}}(\chi, s) = \sum_{\mathfrak{a} \subseteq \mathcal{O} \text{ invertible}} \frac{\chi([\mathfrak{a}])}{\mathfrak{N}(\mathfrak{a})}.$$

One can easily show that this function is absolutely and uniformly convergent for $\Re(s) > 1$ and therefore it defines an analytic function on $\mathbb{C} \smallsetminus \{s \in \mathbb{C} \colon \Re(s) \leq 1\}$. For example, see [26, Proposition 8.1]. We can extend the Hecke $\mathcal{L}$-function to a meromorphic function in the whole complex plane with only one pole for $s = 1$.

One of the most famous and important open problems in mathematics is the *Riemann Hypothesis* and one generalization involves the Hecke $\mathcal{L}$-function. This generalization is called *Generalized Riemann Hyptothesis*, for short GRH, and it predicts the distribution of the zeros of the Hecke $\mathcal{L}$-function. It is known that there is no zero with real part greater than 1 and all zeros with negative real part are the even real numbers $-2, -4, -6, ...$, see [26, Chapter VII]. These zeros are called *trivial zeros* of the Hecke $\mathcal{L}$-function. Hence, all other zeros must lie in the *critical strip*, i.e. the strip in the complex plane

$$\{s \in \mathbb{C} \colon \Re(s) \in [0, 1]\}.$$

The GRH precisely claims that all the nontrivial zeros of the Dedekind $\mathcal{L}$-series lie in the strip

$$\{s \in \mathbb{C} \colon \Re(s) \in [0, 1/2]\},$$

and so the half plane $\Re(s) > 1/2$ is zero-free.

The next result shows that a particular Cayley graph associated with the class group of a quadratic imaginary order satisfies the expansion property.

Theorem 3.4.9. *Let L be a quadratic imaginary number field and $\mathcal{O}$ the order of conductor f in L. Let $\mathfrak{f}$ be the principal ideal in $\mathcal{O}$ generated by the conductor, d_L the discriminant of L and $q = d_L \mathfrak{N}(\mathfrak{f}) = d_L f^2$. Let x be a positive real number. Consider the set*

$$S_x = \{[\mathfrak{p}] : \mathfrak{p} \text{ is an invertible prime } \mathcal{O}\text{-ideal such that } \mathfrak{N}(\mathfrak{p}) < x \text{ or } \mathfrak{N}(\mathfrak{p}^{-1}) < x\}.$$

Then, assuming GRH for the characters of $\mathrm{cl}(\mathcal{O})$*, the graph $G = Cay(\mathrm{cl}(\mathcal{O}), S_x)$ has trivial eigenvalue*

$$\lambda_1 = 2\mathrm{li}(x) + O(\sqrt{x}\log(x|q|)), \quad \textit{where} \quad \mathrm{li}(x) = \int_2^x \frac{dt}{\log t},$$

while the nontrivial eigenvalues satisfy the bound

$$|\lambda| = O(\sqrt{x}\log(x|q|)).$$

If we let $B > 2$ and $x \geq (\log|q|)^B$, then the bound becomes

$$|\lambda| = O((\lambda_1 \log \lambda_1)^{1/2+1/B}). \tag{3.3}$$

Proof. One can find a proof for a more general result in [18, Theorem 1.1]. In order to build a proof for this particular version, one can use [18, Remark 1.2.a] and [10, Theorem 7.22]. □

Corollary 3.4.10. *Let $L, \mathcal{O}$ and q as in the previous theorem and let $x \geq (\log|q|)^B$, for $B > 2$. Let $h(\mathcal{O})$ be the class number of the order $\mathcal{O}$. If GRH holds, then there exists a positive constant C such that, a random walk of length*

$$t \geq \frac{\log(h(\mathcal{O})/\#S^{1/2})}{\log(\lambda_1^{\frac{B-2}{2B}}/C\log\lambda_1^{\frac{B+2}{2B}})}$$

from any starting vertex lands in any fixed subset $S \subseteq \mathrm{cl}(\mathcal{O})$ with probability at least $\dfrac{\#S}{2h(\mathcal{O})}$.

Proof. Following the notation established in the previous theorem, one can just apply Theorem 3.4.5 to $Cay(\mathrm{cl}(\mathcal{O}), S_x)$, using $c = O((\lambda_1\log\lambda_1)^{1/2+1/B}) = C(\lambda_1\log\lambda_1)^{1/2+1/B}$, for some $C > 0$. □

Now that we have these powerful results, we want to apply them to isogeny graphs, starting from the case of complex elliptic curves with complex multiplication.

Theorem 3.4.11. *Let $\mathcal{O}$ be an order of discriminant d and conductor f in a quadratic imaginary number field L of discriminant d_L. Let G be the graph whose vertices are elements of $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ and whose edges are isogenies of prime degree less than $M \geq (\log|d|)^B$, for some constant $B > 2$. Then, assuming GRH, the graph G is a two-sided expander graph satisfying the bound* (3.3).

Proof. By Theorem 2.3.5, $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ and $\mathrm{cl}(\mathcal{O})$ are in bijection. Moreover, isogenies of prime degree less than M correspond to the action of integral ideals of prime norm less than M, and the inverses of such ideals have the same prime norm and therefore they yield such isogenies, too. Thus, the graph G is isomorphic to the Cayley graph of $\mathrm{cl}(\mathcal{O})$ with generating set consisting of ideals of prime norm less than $M \geq (\log|d|)^B$. The desired result comes directly from Theorem 3.4.9 with $x = M$, after noticing that the requirement $x \geq (\log(f^2 d_L))^B$ is satisfied since $d = f^2 d_L$. □

The next theorem performs a reduction and gives an analogous result for ordinary elliptic curves over a finite field, using Deuring's theorems.

Theorem 3.4.12. *Consider the set $S_{N,q}$ of j-invariants of ordinary elliptic curves defined over $\mathbb{F}_q$ having exactly N $\mathbb{F}_q$-rational points. Fix $j \in S_{N,q}$ and let $\mathcal{V}$ be the set of all j-invariants corresponding to curves with the same endomorphism ring as the one of the curve corresponding to j. Form a graph G on the set of vertices $\mathcal{V}$ by connecting j_1 and j_2 if there exists an isogeny of prime degree less than $(\log 4q)^B$ between them, for some fixed $B > 2$. Then, if GRH holds true, G is a two-sided expander graph in the sense that its nontrivial eigenvalues satisfy the bound* (3.3).

Proof. Let $\mathcal{O}$ be the order in the quadratic imaginary number field L such that $\mathrm{End}(E) = \mathcal{O}$ for all curves whose j-invariant is in $\mathcal{V}$. Let d be the discriminant of $\mathcal{O}$. We have that $4q = t^2 - f^2 d$, where f is the conductor of $\mathcal{O}$ and t is the trace of the Frobenius endomorphism. Hence, we have that

$$(\log |d|)^B = \left(\log \frac{|t^2 - 4q|}{f^2}\right)^B \leq \left(\log \frac{4q}{f^2}\right)^B \leq (\log 4q)^B,$$

since the trace t satisfies the Hasse bound $|t| \leq 2\sqrt{q}$. Therefore, $(\log 4q)^B$ satisfies the condition for M in Theorem 3.4.11. Now, we want to show that the graph in Theorem 3.4.11 is isomorphic to the graph G. The curves in $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ are all defined over the ring class field H of $\mathcal{O}$. Identification of vertices is achieved by choosing a prime ideal $\mathfrak{q}$ in the ring of integers of H of norm q and reducing the curves in $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$ to obtain curves in $S_{N,q}$. Theorem 2.4.9 implies that this identification is surjective, so we need to show that it is injective. Consider two non-isomorphic elliptic curves $E_{\mathfrak{a}}$ and $E_{\mathfrak{b}}$ in $\mathrm{Ell}_{\mathcal{O}}(\mathbb{C})$, which means that $\mathfrak{a}$ and $\mathfrak{b}$ must lie in different ideal classes. By Proposition 2.3.16, there exists a prime ideal $\mathfrak{c}$ belonging to the same ideal class of $\mathfrak{a}\mathfrak{b}^{-1}$. Recall that there is only a finite number of ramified ideals, because the norm of such ideals needs to divide the discriminant of the number field L. Therefore, since for each ideal class there are infinitely many prime numbers that are norms of ideals in that class, we can assume without loss of generality that $\mathfrak{c}$ is unramified. In particular, $\mathfrak{c}$ is not a principal ideal, because $\mathfrak{a}$ and $\mathfrak{b}$ do not lie in the same ideal class, and its norm $\mathfrak{N}(\mathfrak{c}) = p$, for some p prime number. This means that $(p) = \mathfrak{c}\bar{\mathfrak{c}}$ and they are the only ideals in $\mathcal{O}$ of norm p. Moreover, since $\mathfrak{c}$ is not principal, also $\bar{\mathfrak{c}}$ is not principal. The ideal $\mathfrak{c}$ induces an isogeny $\phi_{\mathfrak{c}}$ between $E_{\mathfrak{a}}$ and $E_{\mathfrak{b}}$ of degree p. If the reductions modulo $\mathfrak{q}$ $\overline{E}_{\mathfrak{a}}$ and $\overline{E}_{\mathfrak{b}}$ were to be isomorphic, then the reduction $\overline{\phi}$ would be an endomorphism of $\overline{E}_{\mathfrak{a}}$ of degree p. However, there is no element in $\mathcal{O}$ that has norm p. If $\alpha \in \mathcal{O}$ has norm p, then $\mathfrak{N}((\alpha))$ would be p, and so we would have $(\alpha) = \mathfrak{c}$ or $(\alpha) = \bar{\mathfrak{c}}$. This is a contradiction, since they are not principal ideals by construction.

For each prime ℓ, the reduction map sends every ℓ-isogeny in characteristic 0 to an ℓ-isogeny in characteristic p. All the isogenies over a finite field of characteristic p are obtained in this way, since $\overline{E}[\ell]$ cannot have more cyclic subgroups of order ℓ than $E[\ell]$ and ℓ-isogenies are in one-to-one correspondence with such subgroups. □

Remark 3.4.13. The previous theorem relies purely on the correspondence between prime degree isogenies and ideals class of prime norm. Hence, one could also apply the result to the case of $\mathbb{F}_p$-rational isogeny graphs of supersingular curves and prove that also those graphs are expanders, under the right hypothesis.

Proposition 3.4.14. *Let G be the same graph as in the previous theorem, whose vertices are denoted by $\mathcal{V}$. If the GRH holds, then there exists a positive constant C such that a random walk of length*

$$t \geq \frac{\log(h(\mathcal{O})/\#S^{1/2})}{\log(\lambda_1^{\frac{B-2}{2B}}/C \log \lambda_1^{\frac{B+2}{2B}})}$$

from any vertex lands in any fixed subset $S \subseteq \mathcal{V}$ *with probability at least* $\dfrac{\#S}{2h(\mathcal{O})}$*, where* $\mathcal{O}$ *is the endomorphism ring of the curves with* j*-invariant in* G *and* $h(\mathcal{O})$ *is its class number.*

Proof. The proof is obtained directly by applying Corollary 3.4.10 and Theorem 3.4.12. □

Loosely speaking, this last proposition ensures that if we fix a vertex j_1 in the isogeny graph G, we take a long enough random walk from j_1, we store our final vertex j_2 and we forget the path we took, then it is hard to recover the sequence of visited vertices.

Chapter 4

CSIDH

The first proposal for an isogeny-based primitive was made by Couveignes in 1997. His idea yields a key exchange protocol where the space of public keys is the set of isomorphism classes of ordinary elliptic curves over $\mathbb{F}_q$, all sharing the same endomorphism ring $\mathcal{O}$, which is an order in an imaginary quadratic number field. It exploits the $\mathrm{cl}(\mathcal{O})$-action on elliptic curves and strongly relies on the commutativity of $\mathrm{cl}(\mathcal{O})$. Couveignes did not publish any official paper on this topic and, eventually, it was discovered independently by Rostovtsev and Stolbunov in 2004. For this reason, we will refer to this cryptosystem as Couveignes-Rostovtsev-Stolbunov, CRS for short.

CRS has two big flaws. In 2010, Childs, Jao and Soukharev found an attack breaking CRS and a quantum adaptation of this attack runs in subexponential time. As this may still be tolerable, the main concern about CRS is its slowness.

There are two different paths attempted to solve the two aforementioned problems.

The first is to shift our attention to supersingular curves: the attack due to Childs, Jao and Soukharev exploits the commutativity of $\mathrm{cl}(\mathcal{O})$, hence indirectly the commutativity of $\mathcal{O}$. This led Jao and De Feo to consider supersingular elliptic curves, whose full endomorphism ring is an order in a quaternion algebra, which is not commutative. This gave birth to the cryptosystem SIDH, which stands for Supersingular Isogeny Diffie Hellman. However, this cryptosystem shares as public keys some additional data, and this led Castryck and Decru to discover a devastating key-recovery attack, which makes SIDH completely unsecure.

The other available path is to accept the subexponential attack on CRS and try to make the cryptosystem faster. This is the purpose of CSIDH [4], where the C stands for "commutative". While SIDH is pronounced spelling every single letter, CSIDH authors Castryck, Lange, Martindale, Panny and Renes decided that the correct pronunciation is "seaside". The idea is still to shift to supersingular curves, but this time restricting ourselves to prime fields. By Theorem 3.3.2, this constraint implies that the endomorphism ring $\mathrm{End}_p(E)$ of a supersingular curve E is an order in an imaginary quadratic number field, and so we can directly apply the construction of CRS. CSIDH runs over 2000 times faster than the current state-of-art implementation of CRS, which itself presents many speedups and ideas to achieve that speed.

4.1 Why Supersingular?

In order to understand where the main speedup of CSIDH comes from, we need to take a look at the initial parameters for both CRS and CSIDH.

A major speedup for CRS comes from the idea of choosing a field of characteristic p such that $p \equiv -1 \mod \ell$ for all distinct small primes ℓ up to a given bound. Then, one needs to find ordinary elliptic curves $E/\mathbb{F}_p$ such that $\#E(\mathbb{F}_p) \equiv 0 \mod \ell$ for as many primes as possible. Notice that this is equivalent to looking for elliptic curves over $\mathbb{F}_p$ with a point of order ℓ for as many primes as possible. These properties make the computations of the $\mathrm{cl}(\mathcal{O})$-action efficient, as we will see in the following. However, finding an ordinary elliptic curve with those properties is hard, and the main research focus for CRS is on speeding up this search.

In the supersingular case, if we assume $p > 3$, the fact that $\#E(\mathbb{F}_p) = p + 1$ for all supersingular curves implies automatically that $\#E(\mathbb{F}_p) \equiv 0 \mod \ell$ for all primes ℓ dividing $p + 1$. Hence, if we choose $p = \ell_1 \cdots \ell_n - 1$, we can easily satisfy the required condition and choose any supersingular curve.

The choice of supersingular curves boasts many other advantages. The class group $\mathrm{cl}(\mathcal{O})$ is a finite abelian group whose cardinality asymptotically is

$$\#\,\mathrm{cl}(\mathcal{O}) \sim \sqrt{\Delta},$$

where $\Delta = |t^2 - 4p|$ is the discriminant of $\mathcal{O}$, see [31]. Recall that the trace of the Frobenius morphism of a supersingular elliptic curve over $\mathbb{F}_p$ is zero, so the discriminant is as large as possible. Hence, we have

$$\#\,\mathrm{cl}(\mathcal{O}) \sim \sqrt{p}, \tag{4.1}$$

and thus, for a fixed choice of p, the size of the class group is nearly as large as possible. Thanks to this observation, one can choose p to be relatively small, which directly affects the key size positively. This observation combined with Theorem 4.1.1 explains why CSIDH has a really small key size.

As always, assume that K is a field of characteristic different from 2 and 3. The short Weierstrass equation is not the only possible equation for an elliptic curve. In CSIDH, it is preferable to use another one, called Montgomery equation, which is

$$y^2 = x^3 + Ax^2 + x, \tag{4.2}$$

with $A \in K$ and $A^2 - 4 \neq 0$. This condition excludes singular curves. Its j-invariant is defined as

$$j = \frac{256(A^2-3)^3}{A^2-4}.$$

Every Montgomery equation can be converted into a short Weierstrass equation: starting from an elliptic curve E with equation (4.2), we can perform the coordinate change $(x, y) \longmapsto (x - A/3, y)$ and we get a new equation

$$y^2 = x^3 + ax + b, \quad \text{where} \quad a = \frac{3 - A^2}{3}, b = \frac{2A^3 - 9A}{27}.$$

Not all elliptic curves can be put into Montgomery form. However, since we are interested in supersingular curves over a particular type of prime field, we can use the Montgomery form thanks to the following theorem.

Theorem 4.1.1. *Let $p > 3$ be a prime such that $p \equiv 3 \mod 8$ and let $E/\mathbb{F}_p$ be a supersingular elliptic curve. Then $\operatorname{End}_p(E) = \mathbb{Z}[\pi]$ if and only if there exists $A \in \mathbb{F}_p$ such that E is $\mathbb{F}_p$-isomorphic to the curve E_A of Montgomery equation*

$$y^2 = x^3 + Ax^2 + x.$$

Moreover, if such an A exists, then it is unique and it is called Montgomery coefficient.

Proof. See [4, Proposition 8]. □

Notice that if we assume $p \equiv 3 \mod 8$, we have $p \equiv 3 \mod 4$, which implies that the curve of j-invariant 1728 is supersingular. We will refer to this curve as E_0, it has equation

$$y^2 = x^3 + x,$$

and it will be the starting curve for our key-exchange protocol. We would like to exclude the cases in which the curves with j-invariant 0 and 1728 have additional automorphism, as we have discussed in Remark 3.1.3. Since the base field is of the form $\mathbb{F}_p$, we just need to make sure that $\mathbb{F}_p$ does not contain nontrivial elements of order 3 and 4. This happens if the order of $\mathbb{F}_p^*$ is coprime both with 3 and 4. For example, we can see that if we choose $p \equiv 11 \mod 12$, then $p - 1 \equiv 1 \mod 3$ and $p - 1 \equiv 2 \mod 4$. Moreover, notice that choosing $p \equiv 11 \mod 12$ implies that $p \equiv 3 \mod 8$, so that we can apply Theorem 4.1.1.

Now, we present an efficient way to identify supersingular curves over $\mathbb{F}_p$, taken from [4]. This is useful for the key-exchange protocol we will describe in the next section. As we have already pointed out, an elliptic curve over $\mathbb{F}_p$ is supersingular if and only if it has exactly $p + 1$ $\mathbb{F}_p$-rational points. The idea is that if we can find a point in $E(\mathbb{F}_p)$ with order $d > 4\sqrt{p}$ that is a divisor of $p + 1$, we can conclude that E is supersingular. The reason is that the Hasse interval in our case is

$$\mathcal{H}(p) = [p + 1 - 2\sqrt{p}, p + 1 + 2\sqrt{p}].$$

Therefore, if $P \in E(\mathbb{F}_p)$ has order $d > 4\sqrt{p}$, there is only one multiple of d in the Hasse interval. If d divides $p + 1$, we can conclude $\#E(\mathbb{F}_p) = p + 1$.

In the setting for our protocol, we will choose $p = 4 \cdot \ell_1 \cdots \ell_n - 1$, with ℓ_i small odd distinct primes. In this case, a random point on a supersingular elliptic curve over $\mathbb{F}_p$ must have order dividing $p + 1$.

This observation leads to the verification method presented in Algorithm 1. If the condition $d > 4\sqrt{p}$ does not hold at the end of Algorithm 1, the point P has not enough big order. This happens with very low probability, as we have observed. However, one can try to run the algorithm with another random point. In (1), Algorithm 1 gives the correct answer, since $[\ell_i]Q_i \neq O$ implies $[p + 1]P \neq O$, i.e. $\#E(\mathbb{F}_p)$ does not divide $p + 1$. In (2), Algorithm 1 multiplies d (which at the end of the algorithm stores the order of the examined point) by ℓ_i: from the previous command, we deduce that $[p + 1]P = O$; since

$[(p+1)/\ell_1]P \neq O$, ℓ_i must divide the order of P. Hence, if the randomly chosen point has not small order, Algorithm 1 is correct and always classifies properly elliptic curves. There is no possibility of wrongly classifying an ordinary curve as a supersingular one.

The algorithm is very efficient, since it only involves choosing a random point and taking multiples of that point in the elliptic curve.

Moreover, the probability of choosing a random point of large order is very high. Indeed, by [36, Theorem 4.1], we have that

$$E(\mathbb{F}_p) \simeq \mathbb{Z}/n\mathbb{Z} \quad \text{or} \quad E(\mathbb{F}_p) \simeq \mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z},$$

for some integers $n \geq 1$ or $n_1, n_2 \geq 1$, with n_1 dividing n_2. By definition of p, the second possibility cannot occur. Hence, by the Chinese Remainder Theorem and the fact that $E(\mathbb{F}_p) = p+1$, we have that

$$E(\mathbb{F}_p) \simeq \mathbb{Z}/(p+1)\mathbb{Z} \simeq \mathbb{Z}/4\mathbb{Z} \times \prod_{i=1}^{n} \mathbb{Z}/\ell_i\mathbb{Z}.$$

Ignoring the even part, a random $\mathbb{F}_p$-rational point corresponds to a random n-tuple $(m_1, ..., m_n) \in \prod_{i=1}^{n} \mathbb{Z}/\ell_i\mathbb{Z}$. If the i-th entry is nonzero, then the prime ℓ_i divides the order of the point. If we denote by d such an order, we have that ℓ_i divides d with probability $(\ell_i - 1)/\ell_i$, because that is the probability of having the i-th entry nonzero. Hence, heuristically, we can assume that a random point has large order with high probability.

Algorithm 1 Verifying supersingularity

```
Require: An elliptic curve $E/\mathbb{F}_p$, where $p = 4 \cdot \ell_1 \cdots \ell_n - 1$. $O$ is the identity of the group.
Ensure: Supersingular or Ordinary.
  Randomly choose a point $P \in E(\mathbb{F}_p)$.
  $d \leftarrow 1$
  for $i \in \{1, ..., n\}$ do
      $Q_i \leftarrow [(p+1)/\ell_1]P$
      if $[\ell_i]Q_i \neq O$ then return Ordinary          ▷ (1)
      end if
      if $Q_i \neq O$ then $d \leftarrow \ell_i \cdot d$          ▷ (2)
      end if
      if $d > 4\sqrt{p}$ then return Supersingular
      end if
  end for
```

4.2 Key-exchange Protocol

First, we introduce the standard situation in a Diffie-Hellman key-exchange protocol.

We have three characters playing: Alice, Bob and Eve. Alice and Bob need to communicate and share secret information through a channel, and every piece of data passing through this channel can be stored by Eve, the evil character. Eve's goal is to steal Alice's and Bob's secret information. Each user of our system has a set of secret (or private) keys $\mathcal{S}_U$ and a set of public keys $\mathcal{P}_U$. A *key-exchange protocol* is a number of prescribed actions involving the keys $\mathcal{S}_A, \mathcal{S}_B, \mathcal{P}_A, \mathcal{P}_B$ performed by Alice and Bob in order to agree on a common secret key that Eve cannot recover using the information she stored during the process and her own keys $\mathcal{S}_E, \mathcal{P}_E$. We assume that Alice and Bob do not share with anybody their secret keys and that Eve knows how our protocol works.

CSIDH needs the following public global parameters.

(1) A prime number p of the form $p = 4 \cdot \ell_1 \cdots \ell_n - 1$ such that $p \equiv 3 \mod 8$, where $\{\ell_1, ..., \ell_n\}$ is the set of all distinct odd small primes up to a given bound.

(2) The elliptic curve E_0 of equation $y^2 = x^3 + x$ over $\mathbb{F}_p$. Remember it is supersingular, since $p \equiv 3 \mod 4$.

A classical number theoretic result due to Dirichlet ensures that there exist infinitely many primes of the form we require.

Theorem 4.2.1 (Dirichlet)**.** *Let n, m be coprime positive integers. There are infinitely many prime numbers of the form $n + km$, with $k \in \mathbb{Z}$.*

Proof. See [15]. □

The trace of the Frobenius endomorphism π is zero, and so π satisfies $\pi^2 = -p$. By Theorem 3.3.2, $\mathrm{End}_p(E_0) = \mathcal{O}$ is an order in the quadratic imaginary number field $\mathbb{Q}[\pi] = \mathbb{Q}[\sqrt{-p}]$. Precisely, Theorem 4.1.1 shows that $\mathcal{O} = \mathbb{Z}[\pi]$, which has conductor 2.

This choice of parameters implies that the ℓ_i-isogeny graph is a disjoint union of cycles, by Theorem 3.3.8. Indeed, for each ℓ_i, we have that

$$\left(\frac{\Delta}{\ell_i}\right) = \left(\frac{-4p}{\ell_i}\right) = \left(\frac{-p}{\ell_i}\right) = \left(\frac{-4\ell_1 \cdots \ell_n + 1}{\ell_i}\right) = \left(\frac{1}{\ell_i}\right) = 1,$$

where Δ is the discriminant of $\mathbb{Q}[\pi]$. Moreover, since $\pi^2 - 1 \equiv 0 \mod \ell_i$, by the classical number theoretic result [24, Theorem 27], we have that the ideals $\ell_i\mathcal{O}$ split as $\ell_i\mathcal{O} = \mathfrak{l}_i\bar{\mathfrak{l}}_i$, where

$$\mathfrak{l}_i = (\ell_i, \pi - 1) \quad \text{and} \quad \bar{\mathfrak{l}}_i = (\ell_i, \pi + 1),$$

and the length of each cycle equals the order of $\mathfrak{l}$ in $\mathrm{cl}(\mathcal{O})$. We would like to know the exact structure of the class group $\mathrm{cl}(\mathcal{O})$ and be able to sample uniformly elements at random. The cryptographic size of the discriminant makes this currently not feasible. Hence, we need some heuristic arguments. Recall that in the ring of integers of a number field we have unique factorization into prime ideals. However, $\mathcal{O}$ is strictly contained in the ring of integers $\mathbb{Z}[\frac{1+\pi}{2}]$, because $p \equiv 3 \mod 4$. Hence, in general, we cannot expect uniqueness of the factorization in $\mathcal{O}$. We assume that the $\mathfrak{l}_i$ do not have small order in $\mathrm{cl}(\mathcal{O})$ and are "evenly distributed" in the class group, so that we can expect different ideals of the form $\mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_n^{e_n}$ for small e_i to lie in the same class only occasionally. For efficiency reason, it is preferable to sample exponents e_i from a short range around zero. We will show that

$\{-m, ..., m\}$ with m such that $2m + 1 \geq \sqrt[n]{\# \operatorname{cl}(\mathcal{O})} \sim \sqrt[2n]{\Delta} = \sqrt[2n]{4p}$ is enough. As the ideals $\mathfrak{l}_i$ are fixed global parameters, we can represent the ideal

$$\prod_{i=1}^{n} \mathfrak{l}_i^{e_i}$$

simply as a vector $(e_1, ..., e_n)$.

At this point, we need to understand how we can walk on the isogeny graph, i.e. how to compute the target curve of the isogeny induced by the action of $\operatorname{cl}(\mathcal{O})$. We are only interested in the codomain since this is the only necessary information to move along the isogeny graph. Let's say we want to compute the isogeny $\phi_{\mathfrak{l}_i}$, the one produced by the action of the ideal $\mathfrak{l}_i$ on a supersingular curve $E/\mathbb{F}_p$. One method can be the following. We start by finding a basis of the ℓ-torsion subgroup

$$E[\ell] = \frac{\mathbb{Z}}{\ell\mathbb{Z}} \times \frac{\mathbb{Z}}{\ell\mathbb{Z}}.$$

This is a 2-dimensional vector space over $\mathbb{Z}/\ell\mathbb{Z}$. Then, we compute the eigenspace of the Frobenius, i.e. we want to find a point $P \in E$ such that $\pi(P) = \lambda P$. We already know that the eigenvalues are $+1$ and -1. Hence, we may simply choose P among $\mathbb{F}_p$-rational points of order ℓ_i. Notice that the vector subspace generated by such a P is exactly the torsion group $E[\mathfrak{l}_i]$, since we have

$$E[\mathfrak{l}_i] = E[(\ell_i, \pi - 1)] = \{Q \in E \colon [\ell_i]Q = O \text{ and } \pi(Q) = Q\} = E[\ell_i] \cap \ker(\pi - 1).$$

Hence, having computed $E[\mathfrak{l}_i]$, using Velú formulas, we can efficiently compute the target curve of the corresponding isogeny ϕ. By Theorem 2.3.19, ϕ coincides up to isomorphism with the isogeny we were looking for, i.e. $\phi = \phi_{\mathfrak{l}_i}$. If we wanted to compute the action of the ideal $\overline{\mathfrak{l}_i}$, we could go through the same steps, but choose P to be a $\mathbb{F}_{p^2}$-rational point of order ℓ_i that is not $\mathbb{F}_p$-rational. Recall that every supersingular elliptic curve is defined at most over $\mathbb{F}_{p^2}$.

Notice that we can use the Montgomery coefficients to identify classes of supersingular elliptic curves up to $\mathbb{F}_p$-isomorphism, thanks to Theorem 4.1.1. Recall that each j-invariant of a supersingular elliptic curve corresponds to two distinct $\mathbb{F}_p$-isomorphism classes, by Lemma 3.3.3. This is the reason why it is preferable to use the Montgomery coefficient to represent each class. Indeed, if the classic j-invariant were to be used, we would not be able to distinguish between the two different $\mathbb{F}_p$-isomorphism classes.

We are ready to describe precisely how the CSIDH key-exchange protocol works. Each user has one private key and one public key. The private key is a n-tuple $(e_1, ..., e_n)$ of integers, each sampled from the set $\{-m, ..., m\}$. These integers represent one ideal class $[\mathfrak{a}]$ in the way we have already explained. The public key is the Montgomery coefficient $A \in \mathbb{F}_p$ of of the supersingular elliptic curve $[\mathfrak{a}]E_0$, the one obtained by applying the action of $[\mathfrak{a}]$ to E_0.

Now, suppose Alice and Bob have key pairs $([\mathfrak{a}], A)$ and $([\mathfrak{b}], B)$. If Alice wants to agree on a common secret key with Bob, she needs to verify that Bob's public key B corresponds to a supersingular curve over $\mathbb{F}_p$ with endomorphism ring $\mathcal{O} = \mathbb{Z}[\pi]$. We will denote by E_B Bob's curve. This verification can be performed efficiently, using Algorithm 1. Then, Alice applies the action of her secret ideal $[\mathfrak{a}]$ to E_B, ending up with the curve $[\mathfrak{a}]E_B = [\mathfrak{a}][\mathfrak{b}]E_0$. Bob should proceed analogously with his own secret ideal $[\mathfrak{b}]$ and Alice's public key A to compute the curve $[\mathfrak{b}][\mathfrak{a}]E_0$. The shared secret is now the Montgomery coefficient S of the common secret curve $[\mathfrak{a}][\mathfrak{b}]E_0 = [\mathfrak{b}][\mathfrak{a}]E_0$, which is the same for Alice and Bob, thanks to the commutativity of $\mathrm{cl}(\mathcal{O})$ and Theorem 2.1.13.

Algorithm 2 Key-exchange protocol

```
Require: Alice's and Bob's secret keys, $[\mathfrak{a}]$ and $[\mathfrak{b}]$ respectively; the public curve $E_0$.
Ensure: Secret key agreement.
  $A \leftarrow$ Montgomery coefficient of $[\mathfrak{a}]\mathrm{E}_0$          ▷ Alice's public key
  $B \leftarrow$ Montgomery coefficient of $[\mathfrak{b}]\mathrm{E}_0$          ▷ Bob's public key
  Alice computes $S_A$, the Montgomery coefficient of $[\mathfrak{a}][\mathfrak{b}]E_0$.
  Bob computes $S_B$, the Montgomery coefficient of $[\mathfrak{b}][\mathfrak{a}]E_0$.
```

4.3 Security Assumptions

In this section, we first want to recall some basic definitions in *time complexity theory*. After that, we will focus on understanding on which security assumptions CSIDH relies.

Definition 4.3.1. Let $\mathcal{A}$ be an algorithm that performs computations involving the integer n of k bits. $\mathcal{A}$ is said to be:

(1) a *polynomial* algorithm if there exists a polynomial $p(x)$ such that the number of bit operations required to complete the algorithm is $O(p(k))$;

(2) a *subexponential* algorithm if the number of bit operations required to complete the algorithm is $O(2^{k^{\varepsilon}})$, for all $\varepsilon > 0$;

(3) an *exponential* algorithm if there exists a polynomial $p(x)$ such that the number of bit operations required to complete the algorithm is $O(2^{p(k)})$.

The algorithm $\mathcal{A}$ is said to be *probabilistic* if there is the possibility that some inputs do not produce the desired output, in the sense that it may be wrong or inconclusive. The definition of probabilistic algorithm is opposed to that of *deterministic* algorithm, for which every input produces the desired output.

An example of a probabilistic algorithm is Algorithm 1. Indeed, one could choose a random point of small order, which would not produce an answer. A probabilistic algorithm is considered good if it produces the desired output with high enough probability. In this sense, Algorithm 1 is a good algorithm. Typically, the trade-off consists of giving up determinism and achieving a lower time complexity.

Let us look at the key-exchange protocol from Eve's point of view. She possesses her private and public keys $([\mathfrak{e}], E)$ and can store the two public keys of Alice and Bob. If she wants to be able to recover the shared secret S, she needs to compute the curve

$[\mathfrak{a}][\mathfrak{b}]E_0$. The only obvious way to do so is to recover either Alice's or Bob's secret key. This problem is considered computationally hard.

Problem 1. Given two supersingular elliptic curves E, E' defined over $\mathbb{F}_p$ with the same $\mathbb{F}_p$-rational endomorphism ring $\mathcal{O}$, find an ideal $\mathfrak{a}$ of $\mathcal{O}$ such that $[\mathfrak{a}]E = E'$.

Moreover, notice that the ideal needs to be represented in such a way that its action on a curve can be evaluated efficiently, for example, using the method described in Section 4.2. For instance, the ideal could be given as a product of prime ideals of small norms.

Actually, the security of the primitive of CSIDH is based on a slightly different hardness assumption. It is conjectured that the CSIDH primitive is an instance of Couveignes' *hard homogeneous space*, which is a finite commutative group action for which some computations can be performed efficiently and others are hard. See [9].

Definition 4.3.2. A *hard homogeneous space* consists of a free and transitive group action $(G, X, \star)$, where G is a finite commutative group and X is some set such that $\#G = \#X$. The following tasks need to be easy.

(1) Given g_1, g_2 decide whether they are elements of G, compute $g_1^{-1}, g_1 g_2$ and decide if $g_1 = g_2$.

(2) Sample a random element from G with (close to) uniform distribution.

(3) Given x, decide if x is an element in X.

(4) Given $x_1, x_2 \in X$, decide if $x_1 = x_2$.

(5) Given $g \in G$ and $x \in X$, compute the action $g \star x$.

The following tasks need to be hard:

(1) Given $x_1, x_2 \in X$, find $g \in G$ such that $g \star x_1 = x_2$.

(2) Given $x_1, x_2, x_3 \in X$ such that $x_2 = g \star x_1$ for some $g \in G$, find x_4 such that $x_4 = g \star x_3$.

By easy task, we mean a computation for which there exists a polynomial-time algorithm (at least probabilistic). By hard task, we mean a computation for which there is no known probabilistic or deterministic polynomial-time algorithm.

Remark 4.3.3. We would like to make the security of CSIDH rely on what we discussed in Chapter 3, Section 4, because everything we proved is adaptable to the isogeny graphs of CSIDH, thanks to Remark 3.4.13. For the cryptosystem CSIDH, we have $p = 4\ell_1 \cdots \ell_n - 1$ and the choice suggested by the authors of the original paper of CSIDH is $n = 74$. This is a compromise to make the system both secure and efficient. If we take $\ell_1, ..., \ell_{73}$ to be the first 73 primes, the first prime number that makes p a prime number is $\ell_{74} = 587$. In order to obtain an expander graph and take advantage of its rapid mixing properties, after fixing $B > 2$, we should consider the isogeny graph $G_{p,P}$, where P is a set containing all primes up to

$$(\log 4p)^B \sim (350)^B,$$

which is way larger than 587. Hence, we cannot be sure whether the CSIDH isogeny graph is actually an expander. However, as pointed out in [18, Section 7.2], it seems that the requirement $B > 2$ is not sharp and $B > 1$ is expected. If this turns out to be true, by performing a slightly different choice on the primes ℓ_i, we could make the CSIDH isogeny graph into an expander graph.

4.4 Classical Security

By *classical security*, we allude to the resistance of a cryptosystem to key-recovery attempts that can run on classical computers.

The most naive approach to attack CSIDH and recover the private key of another user is to perform a search through all possible keys. This method is called *brute force attack*. Recall that a private key for CSIDH consists of a vector of exponents $(e_1, ..., e_n)$ such that each exponent can take value in the range $\{-m, ..., m\}$. Every vector represents the ideal class $[\ell_1^{e_1} \cdots \ell_n^{e_n}]$, but there may be multiple possible representations, since the order we are working with is not the ring of integers. This means that the morphism of groups

$$(e_1, ..., e_n) \longmapsto \prod_{i=1}^{n} [\mathfrak{l}_i^{e_1}]$$

has nontrivial kernel. Even if it is not clear how to explicitly determine the kernel of the above morphism, we argue with an heuristic argument that the number of *short* representatives per ideal class is small. With the adjective short, we mean that the exponents are taken from the set $\{-m, ..., m\}$.

Assume that $\mathrm{cl}(\mathcal{O})$ is *almost cyclic*, in the sense that there exists a very large cyclic subgroup G. Indeed, [7, 9.1] provides a heuristic argument on why this is true with high probability for an arbitrary imaginary quadratic number field. Let us consider the group homomorphism

$$\rho\colon\ \mathrm{cl}(\mathcal{O}) \longrightarrow \mathbb{Z}/N\mathbb{Z},$$

where N is the order of the large cyclic subgroup in $\mathrm{cl}(\mathcal{O})$. This map is the projection from $\mathrm{cl}(\mathcal{O})$ to G composed with the isomorphism $G \simeq \mathbb{Z}/N\mathbb{Z}$. By assumption, ρ is surjective. Let $\alpha_i = \rho([\mathfrak{l}_i])$. Let us suppose that $\alpha_1 = 1$, which can be done without loss of generality if one of the ideals ℓ_i has order N. If this is not the case, we can replace $\mathrm{cl}(\mathcal{O})$ with the subgroup generated by the ideals $\mathfrak{l}_i$ and everything works fine in a completely analogous way. For any fixed $[\mathfrak{a}] \in \mathrm{cl}(\mathcal{O})$, any of short representations $[\mathfrak{a}] = [\mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_n^{e_n}]$ correspond to a solution to the linear congruence

$$e_1 + e_2\alpha_2 + \cdots + e_n\alpha_n \equiv \rho([\mathfrak{a}]) \mod N. \tag{4.3}$$

Therefore, if we count the number of solutions to this linear congruence, we can establish an upper bound for the number of short representations of the ideal class $[\mathfrak{a}]$. Notice that the number of solutions coincides with the number of solutions of the associated homogeneous linear congruence

$$e_1 + e_2\alpha_2 + \cdots + e_n\alpha_n \equiv 0 \mod N. \tag{4.4}$$

Definition 4.4.1. A *lattice of rank m Λ* in $\mathbb{R}^n$, with $m \leq n$, is a discrete subgroup of $\mathbb{R}^n$ such that the $\mathbb{R}$-linear space spanned by its elements has dimension m. Equivalently, it is a free $\mathbb{Z}$-module of rank m. If the vectors forming any basis of Λ have integer coordinates, we say that Λ is an *integer* lattice.

Given Λ a rank-m lattice in $\mathbb{R}^n$ and a basis $\{u_1, ..., u_m\}$, we can define its volume to be

$$\mathrm{vol}(\Lambda) = |\det(u_1, ..., u_m)|.$$

The volume of a lattice can be interpreted as the m-dimensional volume of the parallelepiped singled out by any of its bases.

The solution of the homogeneous congruence (4.4) are exactly the points in the integer lattice Λ spanned by the vectors $v_1, ..., v_n$, where

$$\begin{aligned} v_1 &= (N, 0, 0, ..., 0) \\ v_2 &= (-\alpha_2, 1, 0, ..., 0) \\ v_3 &= (-\alpha_3, 0, 1, ..., 0) \\ &\vdots \\ v_n &= (-\alpha_n, 0, 0, ..., 1). \end{aligned}$$

Lemma 4.4.2 (Gaussian Heuristic)**.** *Let Λ be a rank-n lattice in $\mathbb{R}^n$ and M a measurable subset of $\mathbb{R}^n$. The* Gaussian Heuristic *predicts that the number of points in $\Lambda \cap M$ is roughly*

$$\frac{\mathrm{vol}(M)}{\mathrm{vol}(\Lambda)}.$$

For a more detailed insight on the motivations underlying the Gaussian Heuristic, see [27, Definition 2.8].

Applying the Gaussian Heuristic to our case, we can deduce that the number short representations for the ideal class $[\mathfrak{a}]$ is roughly

$$\frac{\mathrm{vol}([-m, m]^n)}{\mathrm{vol}(\Lambda)} = \frac{(2m)^n}{N}.$$

We assumed that $\mathrm{cl}(\mathcal{O})$ is almost cyclic, so we have

$$\frac{(2m)^n}{N} \approx \frac{(2m)^n}{\mathrm{cl}(\mathcal{O})}.$$

Hence, if we choose m to be as small as possible such that $(2m)^n \geq \mathrm{cl}(\mathcal{O})$, we can hope that the number of short representations of the ideal class $[\mathfrak{a}]$ is small.

Let now $([\mathfrak{a}], A)$ be Alice's pair of keys. If we assume that the number of short representations for each ideal class is small, then we can deduce that a brute-force attack has at least exponential computational complexity. Indeed, for every possible secret key $(e_1, ..., e_n)$, one needs to compute the action of the corresponding ideal class $[\mathfrak{e}]$ on E_0 and verify if the Montgomery coefficient of $[\mathfrak{e}]E_0$ equals A, using the method we described in

Section 4.2. Notice that the number of bits of a secret key is $n(\log m+1)$. The brute-force attack we have just described has computational cost

$$O(\mathbf{V}^n(2m+1)^n) = O(\mathbf{V}^n e^{n\log m}),$$

where $\mathbf{V}$ is the computational cost for the action of the largest among the prime ideals $\mathfrak{l}_i$ using Velú formula. Notice that for each candidate tuple, it is preferable to apply n times Velú formula for the action of a prime ideal of small norm rather than apply Velú formula only once to a large norm ideal. The computational cost is at least $O(e^{n\log m})$, which is exponential in the length of the private key. Using the estimate 4.1 and our choice or m and n, we can give the computational cost in terms of p: the computational cost of the brute force attack on CSIDH is $O(p)$.

Currently, there is no known classical algorithm running fast enough to pose a threat to CSIDH. Moreover, it is crucial to notice that Shor's algorithm does not apply to CSIDH trivially, because there is no clear way to translate the CSIDH primitive into a discrete logarithm or a factoring problem. Precisely, the reason why the CSIDH primitive is not based on discrete logarithm is that the set of vertices of the isogeny graph we work with, which is $\mathrm{Ell}_{\mathbb{Z}[\pi]}(\mathbb{F}_p)$, does not have a group structure. However, this does not directly imply that CSIDH is quantum-resistant. Indeed, [21] describes a quantum key-recovery attack that runs in sub-exponential time. At the moment, there is no known attack on CSIDH that runs in polynomial time.

Bibliography

[1] T. M. Apostol et al., *Modular functions and dirichlet series in number theory.*

[2] E. Bach, *Explicit bounds for primality testing and related problems*, Mathematics of Computation, 55 (1990), pp. 355–380.

[3] R. Bröker, *Constructing elliptic curves of prescribed order*, Leiden University, 2006.

[4] W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes, *Csidh: an efficient post-quantum commutative group action*, in Advances in Cryptology–ASIACRYPT 2018: 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2–6, 2018, Proceedings, Part III 24, Springer, 2018, pp. 395–427.

[5] J.-J. Chi-Domínguez, A. Esser, S. Kunzweiler, and A. May, *Low memory attacks on small key csidh*, Cryptology ePrint Archive, (2023).

[6] A. Childs, D. Jao, and V. Soukharev, *Constructing elliptic curve isogenies in quantum subexponential time*, Journal of Mathematical Cryptology, 8 (2014), pp. 1–29.

[7] H. Cohen and H. W. Lenstra Jr, *Heuristics on class groups of number fields*, in Number Theory Noordwijkerhout 1983: Proceedings of the Journées Arithmétiques held at Noordwijkerhout, The Netherlands July 11–15, 1983, Springer, 2006, pp. 33–62.

[8] K. Conrad, *The conductor ideal of an order.*

[9] J.-M. Couveignes, *Hard homogeneous spaces*, Cryptology ePrint Archive, (2006).

[10] D. A. Cox, *Primes of the Form x^2+ny^2: Fermat, Class Field Theory, and Complex Multiplication. with Solutions*, vol. 387, American Mathematical Soc., 2022.

[11] G. P. Davidoff, P. Sarnak, and A. Valette, *Elementary number theory, group theory, and Ramanujan graphs*, vol. 55, Cambridge university press Cambridge, 2003.

[12] L. De Feo, *Mathematics of isogeny based cryptography*, arXiv preprint arXiv:1711.04062, 12 (2017).

[13] C. Delfs and S. D. Galbraith, *Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$*, Designs, Codes and Cryptography, 78 (2016), pp. 425–440.

[14] M. Deuring, *Die typen der multiplikatorenringe elliptischer funktionenkörper: G. herglotz zum 60. geburtstag gewidmet*, in Abhandlungen aus dem mathematischen Seminar der Universität Hamburg, vol. 14, Springer, 1941, pp. 197–272.

[15] P. G. L. Dirichlet, *There are infinitely many prime numbers in all arithmetic progressions with first term and difference coprime*, arXiv preprint arXiv:0808.1408, (2008).

[16] R. Hartshorne, *Algebraic geometry*, vol. 52, Springer Science & Business Media, 2013.

[17] H. Iwaniec and E. Kowalski, *Analytic number theory*, vol. 53, American Mathematical Soc., 2021.

[18] D. Jao, S. D. Miller, and R. Venkatesan, *Expander graphs based on grh with an application to elliptic curve cryptography*, Journal of Number Theory, 129 (2009), pp. 1491–1504.

[19] N. Koblitz, *A course in number theory and cryptography*, vol. 114, Springer Science & Business Media, 1994.

[20] D. R. Kohel, *Endomorphism rings of elliptic curves over finite fields*, University of California, Berkeley, 1996.

[21] G. Kuperberg, *A subexponential-time quantum algorithm for the dihedral hidden subgroup problem*, SIAM Journal on Computing, 35 (2005), pp. 170–188.

[22] S. Lang, *Elliptic functions*, Springer, 1987.

[23] J. Li, S. Li, and Y. Ouyang, *Factorization of hilbert class polynomials over prime fields*, arXiv preprint arXiv:2108.00168, (2021).

[24] D. A. Marcus and E. Sacco, *Number fields*, vol. 1995, Springer, 1977.

[25] D. Martin and L. Ahlfors, *Complex analysis*, McGraw-Hill, New York, 1966.

[26] J. Neukirch, *Algebraic number theory*, vol. 322, Springer Science & Business Media, 2013.

[27] P. Q. Nguyen and B. Vallée, *The LLL algorithm*, Springer, 2010.

[28] A. Pizer, *Ramanujan graphs*, AMS IP STUDIES IN ADVANCED MATHEMATICS, 7 (1998), pp. 159–178.

[29] H.-G. Rück, *A note on elliptic curves over finite fields*, Mathematics of Computation, 49 (1987), pp. 301–304.

[30] R. Schoof, *Nonsingular plane cubic curves over finite fields*, Journal of combinatorial theory, Series A, 46 (1987), pp. 183–211.

[31] C. Siegel, *Über die classenzahl quadratischer zahlkörper*, Acta Arithmetica, 1 (1935), pp. 83–86.

[32] J. H. Silverman, *Advanced topics in the arithmetic of elliptic curves*, vol. 151, Springer Science & Business Media, 1994.

[33] ——, *The arithmetic of elliptic curves*, vol. 106, Springer, 2009.

[34] J. Tate, *Endomorphisms of abelian varieties over finite fields*, Inventiones mathematicae, 2 (1966), pp. 134–144.

[35] J. Vélu, *Isogénies entre courbes elliptiques*, Comptes-Rendus de l'Académie des Sciences, 273 (1971), pp. 238–241.

[36] L. C. Washington, *Elliptic curves: number theory and cryptography*, Chapman and Hall/CRC, 2008.

[37] W. C. Waterhouse, *Abelian varieties over finite fields*, in Annales scientifiques de l'École normale supérieure, vol. 2, 1969, pp. 521–560.